

Published and Copyright (c) 1999
All Rights Reserved

Atari Online News, Etc.
A-ONE Online Magazine
Dana P. Jacobson, Publisher/Managing Editor
Joseph Mirando, Managing Editor

Atari Online News, Etc. Staff

Dana P. Jacobson -- Editor
Joe Mirando -- "People Are Talking"
Michael Burkley -- "Unabashed Atariophile"
Albert Dayes -- CC: Classic Chips

With Contributions by:

Bengy Collins
Kevin Savetz
Carl Forhan
Sam Ismail

To subscribe to A-ONE, send a message to: dpj@delphi.com
and your address will be added to the distribution list.
To unsubscribe from A-ONE, send the following: Unsubscribe A-ONE
Please make sure that you include the same address that you used to
subscribed from.

To download A-ONE, set your browser bookmarks to one of the
following sites:

<http://people.delphi.com/dpj/a-one.htm>
<http://www.icwhen.com>
<http://almag.atari.org>
<http://homestead.dejanews.com/ssag>

Visit the Atari Advantage Forum on Delphi!
<http://forums.delphi.com/m/main.asp?sigdir=atari>

=~::~~::~=

A-ONE #0118

07/02/99

~ People Are Talking! ~ CC: Classic Chips! ~ MS' Game Blocker
~ Computer Dungeon Gone ~ AOL To Release 5.0 ~ Core Design News
~ Best Catalog Addendum ~ Interview: Milan News! ~ MGC '99 Update!
~ Compaq: Bleak Future? ~ Sega Bets Everything! ~ bleem Goes Retail

-* HI Releases Atari History CD *-
-* The "Anatomy" of Hard Drive Design! *-
-* SEC Probing Microsoft Accounting Practices *-

=~==~==

->From the Editor's Keyboard
"*****"

"Saying it like it is!"

Well, the vacation is almost over. It's nice to end it with a long holiday weekend, however. The 4th of July is almost upon us already! Where did the past few months go? Oh well... More barbecues planned for this weekend, with friends and relatives dropping by. It's going to be a HOT weekend, so I'm sure the pool will get a workout!

Lots of news this week - probably our largest issue to-date! Perhaps being home for two weeks played a part, but we've also had some large interviews completed that we're including this week. Albert Dayes returns with his CC: Classic Chips column with a terrific interview and article pertaining to "The Anatomy of a Hard Drive". I'm sure that you'll find it as fascinating as I did!

Lots of Atari-related news, gaming news, and general computing world news for you this week. As has been the pattern the past few months, Microsoft continues to find itself under the microscope. This time they're being investigated by the SEC!

Well, I still have a number of things to finish up in preparation for this weekend's festivities so I won't dawdle. Let me just remind you all to be careful this weekend. Use fireworks safely. And most importantly, if you're going to drink, do not drive!

Until next time...

CORE DESIGN - PRESS RELEASE
17th June 1999

NEW F-MAX PRODUCTS being evaluated!

DESTTINY

What is your destiny? If you own an Atari TT then it's unlimited!

DesTTiny allows Falcon expansion cards to be interfaced to the Atari TT. This opens up several possibilities, including the chance to fit Eclipse and Tempest!

PCI EXPANSION

Similar to DesTTiny except this board plugs into a standard PCI slot on PC, Milan, Hades, etc., allowing standard Falcon expansion cards to be fitted.

Now Tempest is available to even more Atari users!

ECLIPSE NOW IN PRODUCTION

Reserve one now to avoid disappointment! Shipping with 4MB ATI RageII graphic card. Multiple PCI expansion being evaluated!
Release Date: 30th June 1999

TEMPEST NEWS

PCBs are currently being finalized.

- * 200MHz part will probably be fitted.
- * Current estimate of speed indicate it will be equivalent to 50-66MHz 68060 when running the 68K emulation software.

We are currently evaluating;

- * USB ports for multiple peripheral expansion
- * Dual-processor support with the Falcon 68030.
- May also work with other accelerator cards, such as AB040!

CORE DESIGN PC SYSTEMS, from only #299.00!

HD DRIVER V7.7 - NOW AVAILABLE! NEW PRICE: #15.95.

The hash sign (#) indicates price in UK pounds sterling

** ECLIPSE - PCI ADAPTOR FOR THE FALCON

PLEASE NOTE:

Eclipse will be shipping on 30th June 1999. Orders now being taken.
If you have already asked for one to be reserved, please confirm this with payment.

Price:

UK - 199.00 pounds

Estimated prices elsewhere:

Euro* - Eu 298

Germany*- DM 598

France* - FF 1990

USA* - \$320

- * Price is dependent on currency exchange rates and import duties.

Distributors:

Germany: AG Computertechnik (also Austria and Switzerland)

e-mail: AG_Comtech@t-online.de Tel: +49 (8331) 86373

France: ACS Production (also Belgium)

e-mail: alainc@imaginet.fr Tel: +33

USA/Canada: to be confirmed

All other countries will be supplied directly by Core Design.

OVERVIEW

ECLIPSE is an adaptor that accepts standard PCI cards. We will be shipping Eclipse with the 4MB ATI range of graphic cards, so Eclipse will offer up to a massive 1600x1200, 16-bit colour or 1280x1024, 24-bit with high speed and refresh rates.

Eclipse offers the best performance of any graphic adaptor, and offers full

access to 14MB upgrades (unlike other graphic adaptors!).

Multiple PCI expansion being evaluated. Price expected to be around #29.00.
Network cards from #20.00!

fVDI

fVDI (Fenix VDI) is part of an on-going development of the Fenix OS, a replacement operating system. fVDI is being used to drive Eclipse and beta-test versions can now be downloaded for testing against a variety of different applications.

<ftp://rand.thn.htu.se/pub/fVDI/BETA/>

NOTE: fVDI also works on standard Falcons.

BENCHMARKS

Currently Eclipse is producing Gembench results with an average of;

1350% *without* NVDI
864% *with* NVDI

See our web site for representative Benchmarks.

PRICE

Complete with fVDI driver software AND 4MB graphic card - only 199.00 UK pounds.

AVAILABILITY

Production models shipping on 30th June 1999.

** TEMPEST

TEMPEST: The FASTEST accelerator for the Falcon - NOTHING can compete with this! A simple, inexpensive way to accelerate a _standard_ Falcon to achieve amazing performance!

--- Now with the PowerPC (PPC) Processor ---

This means Tempest could run with a staggering 266MHz processor! Due to the lack of board space, the proposed AGP slot will not be fitted.

The current estimate of speed indicates Tempest will be equivalent to a 50-66MHz 68060 when running the 68K emulation software.

Our web site now contains an FAQ section to answer your queries.

Currently being evaluated are the inclusion of;

- * USB ports - multiple peripheral expansion!
- * Dual-processor support with the Falcon 030!
- May also work with Afterburner040 - to be confirmed!

AVAILABILITY:

Not yet determined, but hopefully prototypes will be ready by Q2 1999.

A Sad Day...

From: ComDungeon <comdungeon@aol.com>

June 1st came and went and The Computer Dungeon closed its doors for the last time.

It was a great time running the Dungeon; the best experience of my life. But eventually reality set in and there wasn't enough sales or interest to keep the business afloat. I can't begin to relate the feelings I experienced when I decided to close the business and began the process of doing so.

But the saddest part of all was watching over 700 ST & Falcon software titles loaded into the back of the county refuse truck and hearing the sound of the crusher doing its job.

We tried to find the software a good home: we ran a "Garage Sale" during the last two months we were open where people could name their own buying price but only 9 offers were made & 33 titles sold in over 60 days. We said we would give it away FREE to any user group that wanted it and all they had to do was send us a letter on club stationery asking for the software and offering to pay for the shipping charges. But not one letter came.

Over 700 ST & Falcon software titles: from Calamus to Pagestream; from Dungeon Master to Populous; from Bentley Bear to Steinberg. Plus joysticks, mice, cables, two ST's, 8bit software & hardware.... gone.

It was great running the Dungeon.... and sad to see it close.

Thanks to all who supported us.
Al Horton

Best Electronics Catalog Addendum

From: Kevin Savetz <savetz@northcoast.com>

[This news item is a little late, but better late than never...]

Atari dealer Best Electronics has posted an update to its revision 10 catalog.

The update includes new prices, new products and information about out-of-stock items.

Their Web site is at: <http://www.best-electronics-ca.com>

The catalog update is at: <http://www.best-electronics-ca.com/addendum.htm>

Milan Interview: Very Good News

MO: I am sorry to begin with such a difficult question, but please answer as best as possible. In a previous interview with Alexander Clauss, I questioned him about his opinions regarding the Milan computer. He stated that he would not buy a Milan because, for one reason, it offers the same incompatibilities as , for example, MagiC Mac. So I ask you to defend yourself. Why should us Atari user's buy a Milan and run MagiC or N.AES on it as opposed to buying a cheap PC or Mac and run MagiC Mac or MagiC PC?

Goukassian :First of all I'd like to thank you for this interview and for giving us the chance to talk about the present and the future of Milan computer systems.

But now the answer to your first question, which is not very hard to be answered for an Atari-fan. From our point of view everyone who believes in the practicability of Atari-Systems should buy a Milan-System. For sure a Milan is not completely compatible with an Atari ST, but the Atari Falcon (when it was released) was not, too.

We found out a way to make it as compatible as possible, because the Milan-board has two programmable logic-chips. Each time Atari-software is trying to work with a non existing hardware, the Milan-Systems starts an exception and the missing Atari-Chips are being emulated by our logic-chips.

That makes a Milan for sure more compatible than any MAC or PC, and also we have the possibility to continue development on the chip-programming, what means that the Milan's compatibility can be optimized continuously without the need to change the motherboard-hardware.

The reason to buy a Milan would be that people want to have true Atari-feeling based on a modern hardware. Being emulated, the TOS is always based on Windows or MacOS, it always depends on the stability of the based OS.

The Milan TOS is being booted directly out of a TOS-Chip like everyone knows it from typical Atari-Systems, what means that booting only takes a view seconds and you have a pure Atari-Computer.

MO: Many people are wondering how many Milan's are exactly out there. Have Milan Sales been 1)very high 2)average or 3)disappointing?

Could you please tell us an estimate of the current amount of Milan's in circulation?

Goukassian: Milan sales have been good. They may have been even better, if we could have produced as many as we needed. Please understand, but I can not tell you any definitive amount. It's been sold since August 98 and we were able to produce at least 50 per month ...

MO: MagiC Milan has recently been released. In what way do you believe this will change the Milan? Could MagiC Milan possibly replace the slightly modified version of N.AES as the pre-installed OS?

Goukassian: MagiC Milan is a big win for Milan-Users as well as everyone who likes to use MagiC. It looks great, works fine and as you can see on our webpage's benchmarks, it's nearly as fast as the Single-TOS. It makes the whole system look even more modern and friendly... At the moment I

can't tell anything about the future of the pre-installed OS.

MO: A common complaint of MagiC users and MiNT advocates is that MagiC seems to be unstable on certain, if not all, systems (ie. Falcons)
How does MagiC Milan perform? Is it stable? Have you experienced frequent crashes? Furthermore, what are your first impressions regarding MagiC Milan?

Goukassian: As I said: We've been running MagiC Milan for more than one week continuously and it was working very fine. There are not more crashes than with MagiC on any other Atari, and we only do have crashes if the user runs very old or incompatible software like games ...

MO: MagiC Milan does not, like the original, totally replace the Milan OS. In fact, MagiC Milan works together in harmony with it. What is the reasoning behind this?

Goukassian: The reason for this is very easy to describe. As MagiC is a completely new OS that means that the hardware-dependable parts of the OS have to be reprogrammed for any new Atari-System. That's the reason why development for Falcon or HADES eg. took so long, and why the HADES-version is not running very well. But our developers have been reprogramming the hardware-dependable parts of the original TOS, what makes them PCI-, ISA- ... compatible. To make MagiC Milan as compatible as possible we decided to only develop the non-hardware-dependable parts of MagiC. The big advantage is that each time we release new hardware versions MagiC will run with the new also offered TOS-base and never lose its compatibility.

MO: Does MagiC Milan have the same "snappy feel" to it that it does when ran on a Falcon or ST?

Goukassian :I would say it's even better because it's incredibly fast. The subjective feeling is as it is on a Pentium 300 MHz-PC.

MO: The Milan has been great for the Atari world. Although it is debatable if in fact the Milan actually converted PC, Macintosh, or Amiga users, there is no question that it has prevented many from leaving our beloved platform.

What should we expect from Milan Computers in the future? Will we ever see the Milan take advantage of technologies such as PowerPC (or even multiple 68060's) processors, built in USB, Firewire, or even DVD rom support?

Goukassian: YES YOU WILL! We've been quiet for a while but we have the most incredible news for everyone you could imagine. If everything should work fine, in late autumn/early winter we will have a brand new Milan, called "ATARI Milan" that will have SDRAM, USB, an 060-processor, DVD-Rom support, a big range of (new) software ...

This all will be produced by Motorola and AXRO GmbH in cooperation with Motorola Computersystems and it will be sold in several hundred (computer)-shops (!). This all has been top secret for many month, but there is a range of software developers working on Milan-Software, Games, Hardware-Add-Ons, new graphic card-drivers ...

But the will be possible, if we should achieve to get the name "ATARI" from HASBRO in license. So our brand new Milan 060, probably called "Milan 500", could be the one and only Atari-Successor! Good news, isn't it?

MO: For a "dead" platform, the Milan040/060 really has quite a bit of competition! Currently available there is the Centurbo II which offers incredible Falcon performance as well as maintaining compatibility. There is also the Hades040/060, and soon, if we are lucky, the Tempest accelerator (which claims that it will offer PowerPC processor's with speeds well over 200mhz..)

To make things even harder for you, in the near future the Phenix060 is supposed to make an entrance. What is your opinion on the CTII, Tempest, Hades, and Phenix? Does Milan computer welcome a bit of friendly competition or are you guys beginning to feel the heat?

Goukassian: Of course we welcome that bit of friendly competition, but as you can see we're trying to deal in another market. We've been working very long and very hard on that deal, but wait until you see the price (will be published August '99), then objectively I would say that no one will ever pay 3000,- DM for a 040/060-system and 700,- to 800,- DM for a hardware-add-on. And with the new Falcon-compatibility-card for about 400,- DM more, any Milan will become a very powerful Falcon-compatible workstation.

MO: When the Milan arrived, there was already the very powerful Hades040/060 on the market. What motivated you guys to introduce a whole new clone, and how long was the Milan in development?

Goukassian: We developed about one year and our motivation was at least, to make a very much more flexible system that costs less than half as much as the HADES does. I think we succeeded in all points.

MO: It sometimes seems that the German Atari market is blind to the needs of the Atari users who are still left in North America. Where can Americans/Canadians purchase a Milan from?

Goukassian: Yes, it seems so. Regarding Milan-Computersystems we do have a lack of manpower for the distribution in North America. But with the new system we'll have a very strong and internationally working distribution partner. We expect Milan-Systems to be available in Europe in early 2000 and in America from april 2000.

MO: Do you believe the Milan is the future of the Atari Platform? If not, what is? What is your feelings regarding MagiC mac and MagiC PC?

Goukassian: Of course I do believe that Milan is the future of Atari-Systems.

Nothing else can be, because the Milan will be the only ATARI-Clone which will be an original Atari, which will be produced by a big partner like Motorola and which will be built in bigger amounts. I like MagiC Mac and MagiC PC as well as some other emulations, but as I introduced my answers: They are all based on an existing, fat and not always very stable running OS.

MO: The Milan Website is huge, packed with information and downloads, updated regularly, and is available in many languages. There is no question that it is truly one of the most useful and impressive Atari related sites available to users of the internet. Since the arrival of MagiC Milan, will we see a separate "MagiC Section" of the page, containing MagiC Milan related downloads/patches, reviews, and information?

Goukassian: There will be a separate MagiC-Section but not before our

webmaster does have a Milan-Computer. Until now we had to sell any machine we could build, but from June/July on, Joachim Fornallaz, the webmaster, will continue web-site development on a Milan running MagiC Milan.

MO: I noticed that you have a banner on the main page stating "Over 300 000 visits in the first 2 months." If this is true, it is very impressive, and proves once again that the Atari Platform is far from dead. How many visits, on average, does the Milan page receive lately?

Goukassian: To be honest I have not checked the number of accesses for more than 6 months but I can truly assure you that the banner is not a lie. I will try to find out the actual number of website-accesses as soon as possible.

(Note: The Milan Webmasters tells me that the Milan homepage receives app. 200 daily hits)

MO: Now some easy questions. What computers, other then Milans, are used in the production of the Milan and the Milan website? Which software is used by the team on a regular basis?

Goukassian: Actually we use a Mac for the website-development, running MagiC Mac, but only as long as we have a Milan-System left for the webmaster. We also work with PCs for the billing-system, but since Milan060-card is available, we work with Calamus on an Atari-System again. Before it was a Mac.

MO: Why did you stay with the Atari platform for so long, and why do you continue to support it?

Goukassian: We love it and we love the way people love it!

MO: Do you consider the Milan a success?

Goukassian: Yes, I would say so, and especially regarding the future it will even become a bigger success! But we have to keep in mind: All this depends on the Atari-name-license being given by HASBRO. If they should not agree, we'll have problems to launch the system the way we want to.

MO: Finally, is there anything you would like to say in conclusion of this interview?

Goukassian: Yes, we - the Milan-Team - want to thank you all for supporting Atari-Systems for such a long time, and we will try to give back what Atari was obviously not willing to give you back anymore after 1993: We will produce and offer new powerful systems for a low price and with the feeling that they have something very special PCs and Macs do not have!

MGC99 Update

From: "Bengy Collins" <collins@bulli.com>

Hello all,

Here is an update for all who is interested in the MGC99 contest. Currently we have approx. 7 entries. These daring programmers are making a

wide range of GEM games. Some of the more interesting games being made are:

GEM Pipemania
GEM Kicker
GEM Galaga
GEM Asteroids

And the sponsors are coming along nicely. So far we have over 500\$ in sponsor prizes and we plan on adding a lot more! It is almost certain now that every Entry will get some nice prizes.

So far, our sponsors are:

AtariComputing Magazine who will donate free mags/renegade CD's to all entries.
M.u.C.S who has donated over 190\$ in software
R.O.M Logicware who is giving one lucky contestant a free copy of the latest Papyrus Office. (app. 150\$)
6 Radium CDs [Radium CD's include tracks from dexter.d and m.pire (produced with Ataris), raytraced animations, a comic, paintings, starmaps and more. Each CD is worth about 22\$]

So for all Atari users interested in Games.. whether you use a Hades, Milan, Falcon, ST, MagiCPC whatever, there should be a much wider selection for you to choose from soon.

(PS for people wanting to enter the contest, you still have plenty of time)

MGC99 English page 1: <http://mgc.atari.org>
MGC99 English page 2: <http://www.atari-computer.de/xonline/mgc/contest.htm>

Contact: collins@bulli.com or thex@atari.org

Thanks!

--

Bengy Collins collins@bulli.com

Just a quick follow up to this:

We learned today that Woller kindly is providing us with *3* copies of the highly anticipated N.AES Version 2.0.

This means that the first place, second place, AND third place winners will all also receive N.AES 2.0.

VINTAGE COMPUTER FESTIVAL - VCF 3.0

From: Sam Ismail <siconic@jasmine.psyber.com>

Parallax Inc. and Dr. Dobb's Journal proudly present...
The Vintage Computer Festival 3.0
October 2nd & 3rd
Santa Clara Convention Center
Santa Clara, California

For the third year in a row, the Vintage Computer Festival is back to thrill and amaze you with the history of one of the greatest inventions in all of mankind, the computer!

This year:

- * More great speakers!
- * More great exhibits!
- * More great vendors with Cool Stuff to sell!
- * More fun games and contests, including the Nerd Trivia Challenge!
- * Tour The Computer Museum History Center!
- * And of course, FREE STUFF!!

In short, plenty of stuff to indulge all your SENSES:

HEAR the speakers! SEE vintage computers! TOUCH them too!

SMELL their electronics! But please don't TASTE them!

"The mission of the Vintage Computer Festival is to promote the preservation of 'obsolete' computers by allowing attendees to experience the people, stories and technologies that embody the remarkable tale of the computer revolution."

Register Early and \$AVE!

Register by September 15 and pay only \$20 per person. That includes access to the entire event all weekend, including the speakers, the exhibit and the flea market. Kids 17 and under are admitted free, and as always, parking is free too!

To pre-register, send a check or money order to:

Vintage Computer Festival
4275 Rosewood Drive #29-161
Pleasanton, California 94588

Or visit the VCF website for more information and a handy registration form ready to print (<http://www.vintage.org/vcf/register.htm>).

The Speakers

Horst Zuse

Horst, the eldest son of early digital computer pioneer Konrad Zuse, will be giving a lecture on his late father's machines and will focus on the computer businesses that Konrad later founded to market his ideas.

Jon Titus

Who can forget one of the earliest of kit microcomputers, the Mark-8? Jon designed the Mark-8 around Intel's 8008 microprocessor and wrote a construction article which appeared in the July 1974 issue of Radio Electronics.

Todd Fischer

Todd bought the assets of IMSAI after their bankruptcy and continued to manufacture and sell IMSAI products under the Fischer-Frietas Corporation. Todd consulted with MGM for the movie War Games and, as a treat, will be bringing along for show and tell the actual IMSAI 8080 computer that co-starred in the film.

Stan Veit

Stan is an early chronicler of the microcomputer revolution, and also a participant. Stan started one of America's first computer stores in New York, and in the process forged relationships with some of today's revered pioneers, including Steve Jobs and Steve Wozniak, whom he shared a booth with at the first West Coast Computer Faire where the Apple][was debuted to the world.

Jim Willing

Our admired and beloved VCF Celebrity, Jim will wax philosophic on various themes of stupendous relevance.

Eugene Kim

Eugene is a writer for Dr. Dobb's Journal (a VCF sponsor!) and will be giving a lecture on the history of free software, a very relevant topic considering the free-software movement is in full force in today's computing arena as evidenced by the escalating adoption of GNU/Linux.

Curt Vendel

Curt is the founder of the Atari Historical Society (check out his great website at <http://www.atari-history.com>) and has amassed an amazing collection of Atari artifacts and company history, including some one-of-a-kind Atari prototypes. Curt will be bringing along some of these prototypes for show and tell.

Roger Sinasohn

Roger is an avid collector of early portable computers and will be giving a talk on the history of the Kyocera 85, a.k.a. the Tandy Model 100, a.k.a. the Olivetti M10, a.k.a. the NEC PC-8201A.

Doug Salot

Doug's research has uncovered what he believes to be the first "personal computer". Forget the Apple][, the Altair or even the Alto. Doug's pick as first PC goes all the way back to the 1950's. Designed by Edmund Berkeley, an unsung hero of computer history, Doug intends to win Berkeley the prominence and praise that he feels is much deserved by this virtually unknown computer visionary.

Plus!

Workshops by Parallax on their amazing little BASIC Stamp. Get a preview of this cool little device by visiting their website at <http://www.parallaxinc.com>.

More speakers to be announced soon! Check the VCF website for updates.

The Vintage Computer Faire and Exhibition

This year, the Vintage Computer Festival exhibit will feature the exhibits of individual collectors, all vying to win first, second or third place in eleven categories, plus the coveted "Best of Show" award. The winner of the "Best of Show" ribbon will also win fabulous prizes.

This year, we want YOU to exhibit the pride of your collection. What better way to demonstrate your enthusiasm for your hobby than to be an exhibitor at the Vintage Computer Festival.

Exhibitor information can be found on the Vintage Computer Festival

website (<http://www.vintage.org/vcf/exhibit.htm>).

The Vintage Computer Flea Market

Every year the best deals on vintage computer items can be found at the Vintage Computer Flea Market. Here are only a few of the amazing items sold at last year's flea market:

- * Apple 1 (no, we're NOT kidding!)
- * SWTPc 6800
- * A rare computer by "the digital group"
- * Morrow Decision 1 (S-100)
- * California Computer Systems (S-100)
- * Apple Lisa 2
- * And much, much more!!

If you've got old computer stuff that you'd like to sell, there's no better place than the Vintage Computer Flea Market at the Vintage Computer Festival. Please visit our website for rates and regulations (<http://www.vintage.org/vcf/vendor.htm>).

The Nerd Trivia Challenge!

Hey trivia fans! Think you got what it takes to survive the Nerd Trivia Challenge? If you qualify, you'll be pitted against two other players in a battle to determine who is the undisputed master of computer history trivia. Answer the most questions correctly and you could walk away with \$50 in cash plus some other nifty prizes. Lightweights need not apply, this contest is geared towards the serious computer history enthusiast. For contest rules please visit <http://www.vintage.org/vcf/ntc.htm>.

Tour The Computer Museum History Center!

The Computer Museum History Center boasts the largest collection of historic computer artifacts in the world. Their collection includes rare, one-of-a-kind computers and artifacts that date back to the early part of the century and beyond.

A tour of the History Center's "Visible Storage" warehouse will be organized on Sunday, October 3. Tickets are limited and so advanced registration is recommended.

FYI: The History Center is on the move, and is thrilled to announce plans to build a permanent museum to be dedicated by 2007. You can help the History Center achieve its goal by becoming a member. Information about becoming a member of the History Center can be found at <http://www.computer-history.org>.

Meet Fellow Computer Collectors!

The VCF is the premier gathering for collectors of vintage computers. What better way to enjoy your hobby than by meeting other like-minded individuals to share stories and information about your computer collection. Register now! <http://www.vintage.org/vcf/register.htm>

Visit our Sponsors!

Without our sponsors, the Vintage Computer Festival could not be a reality. It is our wonderful sponsors that enabled us to bring

you the Vintage Computer Festival year after year. Demonstrate your appreciation by visiting them on the web:

Parallax, Inc.

<http://www.parallaxinc.com>

Parallax manufactures the BASIC Stamp, a versatile microcontroller with a built-in BASIC language that is suitable for most any electronics and computer hobbyist project.

Dr. Dobb's Journal

<http://www.ddj.com>

Dr. Dobb's Journal is the oldest continually running microcomputer journal, in existence since 1975. It features monthly articles and columns geared towards the computer programming professional.

==~==~==

PEOPLE ARE TALKING
compiled by Joe Mirando
jmirando@portone.com

Hidi ho friends and neighbors. Are you ready for the long weekend? I know that I am. Heck, we work hard all year long. We deserve the measly ten holidays a year that most employers allow us. If it were up to me, we'd have twice that many.

Fortunately, I hold that belief only for people, not for computers. I believe that a computer should be able to handle anything and everything that you throw at it (unless you're trying to get it to do something really stupid, of course).

That's why my TT hasn't been powered down in almost a month, the Stacy sits at the ready, eagerly awaiting my command, and even that new-fangled PC laptop does double duty looking for ET's phone number (remember last week's column?).

I know some people who treat their computers like people... but then again I know others who treat people like computers, so I guess it evens out. Anyway, those who treat their computers like people are really kind of missing the point of having a great tool like a computer at their disposal. Notice that I said "tool"? That's exactly what they are. They are extremely flexible and useful tools to be sure, but tools none the less.

Some of you may remember the go-arounds we had with another editor on an almost weekly basis. What I said then still goes: If your computer does what you need it to do, use it. Use the heck out of it, and then use it some more. Computers don't have feelings (yet) and don't take it personally if you give them one ten minute task after another. But when you can't do what you want or need to with it, go shopping.

My work load has become such that I simply needed a "WinTel" computer. Of course, the first thing I did was to download the free version of

Gemulator 98 and see just how well it works. It does work quite well, by the way. I find it to be a bit quirky, but that might just be me getting used to the new laptop. The speed increases that GemBench reports are simply amazing. And I also took a few minutes to install MagiC (the Atari version) on it. Perhaps one day, if I can find the original disk, I'll install PC Ditto so that I can emulate a PC on an Atari emulator on a PC.... Ouch, my brain hurts. <grin>

Well, enough of this silliness. Let's take a look at what's going on with the UseNet.

From the comp.sys.atari.st Newsgroup

"Fibrillation" asks for help with cross-platform transfers:

"I found something on the web for my Atari and downloaded it to my mac. It's Stuffed in the standard Mac way but it won't allow me to decompress it in order to put on a Dos floppy for transfer to get it to my 140 ST.

Any suggestions?"

Nicholas Bales tells Fib:

"Do you mean is a hqx or sit file ? If that's the case, then I'm 99.999% sure that it is not for consumption on an Atari.

Atari compressed formats are lzh, zip, arc, zoo, and maybe a few others. zip files can be unzipped by MacZip, the others will probably have Mac de-arcers too, but you will usually want to decompress them on the ST. You can get programs for these on most FTP sites. Check out the Quick FAQ."

Martin "Nightowl" Byttebier tells Nick:

"Luckily you didn't say 100%. *.hqx files (almost all) can be unpacked on the Atari. On the Belgian ftp-server you can find a little utility called 'xbin.zip' It's just the binary, there are no docs. To use it just drag & drop the *.hqx file you want to unpack to the program.

ftp://193.190.204.128/atari/archivers/xbin.zip

Oops I didn't read your posting very good. Of course if it's a *.hqx archive it very unlikely it's intend to be used on the Atari."

Hallvard Tangerras tells Fib:

"Decompressing .ZIP files works fine on the Mac (I personally use "Zipit" for this, but there are many others. When it comes to .LZH/.LHA I have problems. There's "MacLHA", and another one I can't remember the name of, but for some idiotic reason it can't understand the MSDOS/Atari syntax of folders, so if you have a folder it gives the file a name such as:

/folder1/folder2/name

Silly! I've tried contacting the authors of these programs, but haven't heard from them. Neither have I had any response in newsgroups. Why not use your Atari for this? Or MagiCMac if you don't have an Atari and want

to use your Mac as an Atari and Mac! It works brilliantly here.

And while I'm here I might as well mention that I performed a little test, and found out that compressing/decompressing worked about twice as fast on the Atari side than on the Mac side with similar Mac software!"

Neil Roughly tells Fib to...

"Use Stuffit Expander. Just make sure the Stuffit Engine is present, which is required for LHa and Arc support. LHa folder structures will remain intact after decompressing. In rare cases an "unsupported" error will occur, whereas MacLHA and LHA Expander will handle them."

Mike Clark asks for help with NVDI:

"I am having problems with NVDI 5 and MagiC 6 on a 4 Meg STe. Whenever I try to load NVDI on startup it locks up the computer"

Steve Sweet tells Mike:

"First off check that all the paths are pointing to valid folders, and if that is so, delete the nvdi.fif file and re-boot, be prepared to wait a short while whilst a new one is created.."

Stephen Barszczak posts:

"I am having something very weird occur with the latest version of Sting in that when it is installed and I am using any software or even just opening windows for any of my drives, the window opens but display a blank window, and the the ST locks up.

Does anyone happen to know why it's doing this?"

Nick Bales tells Steven:

"I must admit I had some window blocking problems with STinG too. I've narrowed it down to having the dialer installed as an accessory. If I use it only as a PRG or APP, it works fine. As an accessory it seems to mess with mouse clicks. This occurred with STinG 1.15 too.

The symptom is usually that at some point, mouse click events are no longer registered, drop down menus work, and you can see the mouse clicking on 3D buttons, but nothing happens and the only way out is to reset."

Ronald Andersson tells Steven:

"STinG itself does not link into any filesystem, and does not directly interfere with any window operations. But it is possible that some mistake in the installation has crashed some memory pointers somewhere, and in that case there is no telling what they might point to, nor how that might interact with other programs.

I hope you did read my UPDATE.TXT files (that is why I wrote them). You can't succeed with updating an old STinG installation if you do not follow the rules of updating as described in those files. Some old parts are incompatible to new parts, and those old parts must be removed, or they will cause some of the new parts to crash.

Installing STinG 1.20 is almost like starting anew, except that you can

keep the old DEFAULT.CFG, CACHE.DNS, and ROUTE.TAB. The STING.PRT file used for STNGPORT.CPX may also be used as basis for a new one, but you must save a new one (as the port drivers have been updated) after checking that port settings are ok with the new STNGPORT.CPX. (All this after installing all stuff and rebooting of course.)

You will also need to reactivate STING.CPX, set its 'delay' to 10 or 20 ms (depending on how fast your machine is), and save those settings.

That (1.08) is a very old version. Did you miss the previous release? (1.15)

As you have been using so old a version, there are many new features in the dialer settings that you should check out. And don't forget that the new dialer has built-in configuration dialogs. You don't have to do all of that with a text-editor anymore (though the SCR files are still ascii).

The MIDI.STX was intended for a future 'ring-type' network with lesser delays than we get simply by using Midi ports connected in a ring today. As yet no one has wanted to take up that project, so the current Midi driver remains (as it always has been) the SERIAL.STX module.

CENTR.STX is another project that has never been finished, although I did some work on it that I have had to abandon due to lack of time for debugging it. When/if it gets completed it will allow a fast link using the printer port, thus allowing STs a fast parallel link

As STinG does not need any 'placeholders' to allow future modules to be installed, I have decided to remove these files from the distribution, until such time as we can provide functional versions.

It is also *vital* that you erase those old modules when installing STinG 1.20, as the port driver structure has been changed to allow a necessary function to be added to all new port drivers. An old module using the old structure will cause STinG to malfunction, so none of those can remain when you install the new STinG 1.20 modules."

Djordje Vukovic tells Ronald:

"I was (I think) quite careful to replace all components of STinG when I installed V1.20; At first everything seemed to work OK: CAB 2.6, Newsie 0.94 in e-mail and ftp functions (I do not know about news- did not have newsgroups available at my ISP at that time).

But then I noticed a weird problem in cooperation of the new STinG and Newsie 0.94: when sending LONG e-mails from Newsie (for example an e-mail with a file of at least 80KB attached), the transmission of e-mail never seemed to end and could not be interrupted either, so I had to reboot.

The problem was repeatable: sending short e-mails was OK, long ones hang the computer. When I deleted STinG V1.20 and reinstalled V1.15 everything worked OK.

I have set the DEFAULT.CFG, DIAL.SCR and the CPXes for STinG V1.20 as I thought should be OK, and after this problem appeared I rechecked that settings were identical as in V1.15; no help. Then I copied the DEFAULT.CFG and DIAL.SCR from V1.15 and used them with V1.20, but with no change.

A detail that might be significant: instead of regular exchanging of "send" and "receive" LEDs on my modem when sending a long e-mail with Newsie+STiNG 1.15, with STiNG1.2 the "send" light is on most of the time, and "receive" blinks only for barely noticeable instants."

"Clint T" asks for help with his ST:

"My question is, Why the HELL is the screen SOO small.. I'm not talking about the monitor itself, I'm talking about when I turn on my ST, the screen is like 1 inch above, 1-1/2 below and 1-inch on both sides inside a box... like it is really 30-40% smaller than the true size of the monitor..

What's the deal with that?

Is the color monitor like that as well? I was thinking about getting one since you can't do anything with out one really...

I also need a hard drive and software! Anyone help me out??!? =)

Would you recommend the ST1040 for a computer, I have it, but I rather have a TT and definitely a FALCON!! (I wish!) anyone wanna do a lot of trading for a falcon?? I want one!

ps, anywhere on the web I can get freeware for the ST for now? or something, I don't even have a hard drive, heh.. but trying to make the most I can out of what I got!

Rich Davey tells Clint:

"I assume it's an original Atari SM monitor? (mono one)

Why not plug it into your TV? Then it'll be as large as you want! and for games/demos and various applications you'll need one anyway (or a colour monitor, which will get a better quality picture, but a TV will suffice, trust me!)

Don't know where (geographically) you are, but consider an IDE Burster kit. Will let you fit a standard PC IDE hard drive to an ST. Doesn't cost the earth either. ASP (in the UK) sell them."

Nicholas Bales replies to Clint:

"Why don't you get a 17" VGA and plug it into that. Don't forget that whatever the size of the screen, resolution will always be 640x400, so having a larger screen will only mean having larger pixels, not a larger workspace.

Software, what kind? There are loads of FTP sites around. For games? <http://lgd.atari.org> contains a few (thousand :)

You should get a Falcon if you want CURRENT and cool software. But hey, the ST is bloody good too, so enjoy!"

Ted Leonard asks for help finding something special:

"I need a macro utility for the Atari ST which has the capability to record mouse clicks, and key strokes, and then allows them to be use when needed. I have used this utility, and believe it is an accessory

(.acc). If you have a copy and could send it to me as an email attachment, or even if you could let me know the name, I could probably find it on the web."

Tom Derrick tells Ted:

I use CodeKeys, which is commercial.

Seems to me that there is a free/shareware one called 'Monkey'. Try a search for that; if no go I'll try to find it and send it to you."

Well folks, that's it for this week. I know that it's a bit short this time around, but I'll make it up to you in the next few weeks somehow or other.

Have a happy, healthy, and safe July 4th celebration, and don't forget to listen to what they are saying when...

PEOPLE ARE TALKING

=~::~~::~=

CC: Classic Chips
= = = = =
by Albert Dayes
70007.3615@compuserve.com

The "Anatomy" of Hard Drive Design
Copyright 1999 - Albert Dayes

Portions Copyright 1997-1999 John Treder

It is a rare opportunity to be able to get an inside view of some the computer technology that we use everyday and this interview is one of those moments. Mr. John Treder has worked at several different companies including IBM, Conner, Priam and Quantum designing hard drives. John has graciously agreed to give us an inside look from the hard drive designer's point of view on how hard drives are created.

Table of Contents

- (1) Birth of a Hard Drive : An Interview with John Treder
- (2) History of Hard Drives
- (3) Drive Performance - TMR
- (4) Seek Time
- (5) Disk Layout - Data Sectors and Servo Sectors
- (6) Glossary of Terms
- (7) References


NOTE:

Mr. John Treder has also previously written more details on hard drive

history, design, and performance issues which are included in sections 2, 3, 4, and 5. A number of edits have been made to the previously mentioned works to clarify and expand on some of the details.

=====
Birth of a Hard Drive : An Interview with John Treder
=====

A VIEW FROM THE SIDE OF A HARD DRIVE:



AD = Albert Dayes
JT = John Treder

AD:

How big is the design team for a hard drive?
What type of members comprise a design team? (e.g. Electrical Engineers, Mechanical Engineers, etc.)

JT:

The size varies, depending on the product's specifications, whether it's a new design or an update to an existing design, and the urgency. The team for Quantum's Viking II could be thought of as more or less typical.

There were half a dozen mechanical engineers, two technicians, seven electrical engineers, one and a half circuit board layout designers, a heads and media engineer, eight people who wrote control and interface code, three product support engineers, two product test engineers, and half a dozen managers. Quantum is unusual in assigning product support engineers near the beginning of a program. It helps a lot, because they contribute to making the design robust from the beginning, and they understand the ins and outs very thoroughly, so their support after the drive is in production is based more on knowledge than guesswork.

In addition to the team members, there are purchasing, marketing, documentation, servowriter development and laboratory specialists who get involved, though they're in "service groups". There can be half a dozen of those people who are "project assigned".

AD:

Is price point the most important factor in a new hard drive?

JT:

Yes. I suspect it even overrides profitability, though hard drive companies can't afford many "loss leaders".

AD:

How long is the development cycle from original design specification to production?

JT:

Again, it varies. A "bump" product, changing the amount of data per platter, can be done in as little as 6 months. A design from scratch, starting a new product family, will take from 2 1/2 to 5 years, but if it takes more than 3 years, something went wrong. Viking 1 was a crash program for Quantum, a new design and new product series, and it took 16 months. Viking II was a significant enhancement of Viking 1, and it took just over 2 years.

AD:

In the late 80's or early 90's one hard drive maker (small company from what I recall) sued all of the other hard drive manufacturers over 3.5 inch disk patents which the company owned. There does not seem to be much litigation in the hard drive business today. Is this due only to the extensive patent cross licensing between the hard drive manufacturers?

JT:

That was Rodime. Conner settled for an amount rumored within the company (I worked there at the time) to be > \$10 million. Quantum fought, and eventually won, but probably spent more money than Conner. I think the major hard drive manufacturers' upper management has recognized that they'll spend less money on the lawyers if they cross-license than if they sue.

Of course, there aren't nearly as many hard drive companies as there were 10 years ago, so there aren't so many opportunities for litigation. The number of manufacturers has decreased enormously--there are only 7 or 8 that I can point to today, where in the mid to late 80's there were at least 50.

AD:

Can you describe the HD design process from your (Mechanical Engineer's) point of view?

JT:

Before a program formally begins, there's a period of product definition. People study head and disk technology, look in their crystal balls (cracked and cloudy, of course) to guess what the hot product is going to be in two years, and generally try to figure out how many disks, what kind of heads, what RPM and what capacity the drive should have. Engineering gets a hand in the process. Usually marketing sets up an approximate specification and the guffaws from engineering cause changes.

There are usually two or three Engineering builds, then two or three Pre-production builds before mass production begins. Before the first E-build, we often take an existing drive and try one or two ideas on it.

We'll only build two or three samples to test the ideas. The first E-build will be 10 or 20 units. The base is either machined from a solid block of aluminum, or modified from an existing base. These units will have the right number of heads and disks, and any key mechanical developments. They'll usually run on a previous product's PCBA, or a new PCBA that's out of form factor and has a socketed processor so the EE's can plug in an emulator (ICE).

Massive changes can happen between the first and second E-builds. Number of disks, RPM, head technology (MR to GMR, for instance), and spindle motor internal design have all changed in my experience. For the second E-build we try to have all the production technologies, though very often the base casting and the actuator won't be made by production methods. At this time, the circuit board will be specific to the product, though it's often still out of form factor and always socketed. The drive will have close to the right TPI, but it usually won't be formatted to full capacity. Data and servo formats and details change almost weekly at this stage.

There will be maybe 100 drives built for E-2. Mechanically, we'll be measuring runouts, shock and vibration performance, EMI (electromagnetic interference), contamination and sealing issues, acoustics, seek performance, and whatever else we can think of. It's our last chance to find and fix major problems without affecting the program schedule.

The first P-build is critical to a project's success. It's when it all comes together. Mechanically, the castings will be castings (not machined from solid), stampings will be stamped, and the drive will generally look like a production unit. Electrically, they'll have real silicon for the main ASICs. For the first time, we'll try to have at least some of the drives written to full capacity. At Quantum, they build 1,000 or 1,500 drives at this time. Formal product testing begins. We show samples to OEM customers, but don't give them any.

The second P-build should be a cleanup of problems found so far, and usually doesn't involve major mechanical changes. If big changes are needed in the mechanics, a third P-build will usually have to happen. P-2 drives are given to OEM customers to begin evaluation. They should show 90% of the production units' performance. Code changes come two or three times a week now.

P-2 is usually the darkest time, emotionally. You have a year or more invested, and all you can see are problems. Marketing has been asking why it wasn't ready 6 months ago.

Then you have the mass production launch. There isn't generally much to do, mechanically, at this stage. If there is, you've got BIG trouble! Servo engineers and interface code guys are working like mad to squeeze the performance and get the bugs out.

AD:

What is the hardest part of the design and/or the design process?

JT:

Inventing a good solution to a key problem. That's also what's the most fun. When you find an answer, it's just wonderful, and usually involves some amount of serendipity. As an example, we were having intermittent spindle motor performance problems with Viking 1, during the E-2 time. Motors that performed poorly also often, but not always, made a

buzzing noise. My boss happened to pick up a motor that had been torn apart and twisted the stator, and it came off in his hands. Now, the stator is supposed to be firmly fastened down!

So, over the next month, the motor company engineers and I worked out several ways to improve the stator's fastening to the base. We ended up adding about 30 cents to the motor's cost, but Viking 1, Viking 2 and Atlas 4 have among the quietest motors in the industry.

AD:

Are the interface/firmware engineers (I would assume EEs mostly) involved from the beginning of the design or only after the mechanical engineers are done?

JT:

The interface guys start a few months after the MEs start. They usually begin to work on some of the E-1 drives and are fully involved by E-2.

AD:

What is the most interesting discovery that you made (which is not a trade secret) during the course of your design or re-design?

JT:

The one that's most interesting to me is maybe a little bit abstruse. The disks and spindle motor bearings and the base casting all form a complicated set of springs that vibrate with various frequencies as the drive spins. The source of some of the vibration is obvious-imbalance in the disks, irregularities in the ball bearings, torque pulses from the motor, for example. But especially at 7200 RPM and above, there's more vibration amplitude, especially at higher frequencies, than these sources ought to produce.

It's air turbulence. At 5400 RPM, most of the air flow is laminar. At 7200 RPM, the airflow at the outside of the disks is in the transition range between laminar and turbulent. And that causes odd vibrations to come and go, as the airflow changes back and forth.

AD:

Do you use computer simulators for all of your designs? (e.g. hardware-software co-design simulator software similar to the software that www.cardtools.com provides)

JT:

The EE and code guys do a lot of simulation. I'm not sure what tools they use. Sometimes it looks a lot like Doom. <vbg> The servo guys use a combination of Matlab and custom-developed software.

Mechanically, we do most of our design these days on a solid modeler. Different companies use different solid modeling software. We also use FEM extensively. One person in each mechanical team is usually proficient with FEM. We also do Monte Carlo analysis of assembly tolerances.

AD:

Can you discuss/define the terms FEM and Monte Carlo analysis?

JT:

FEM = Finite Element Modeling. You use programs such as Nastran or Ansys or Algor or Fluent to model various mechanical problems--stress and deflection, magnetic fields, fluid flow, vibration modes, and so forth.

To become really proficient at using an FEM program requires working at it full time for a year or two. Once you've learned one of them, you can pick up another in six months or so. But I've never met anyone who could use more than one FEM program at the same time--they're all very finicky and different from each other.

You can get results quite easily. It's very difficult to get meaningful results that correlate well with experiment.

FEM used to be mainframe or mini stuff. It started to run on Unix workstations about 10 years ago, and in the last 2 or 3 years it's become practical to run FEM on a high-end Windows NT box. Models run in from a few minutes to overnight. If you have a really big, slow model, it could have taken months to build, and a weekend run doesn't seem all that slow. Big models, in unix, can usually be parceled out across various machines on your local network, over night or over the weekend. FEM is basically inverting a few xillion enormous matrices. The problems involve ill-conditioning and slow inner loops. That's why it takes an expert to get good results--ill-conditioning, especially, can give very bad answers without being obvious.

Monte Carlo analysis is used for studying the effect of assembly variables. "Monte Carlo" refers to "rolling the dice". Any time you have many statistically independent variables, for each one of which you can propose a statistical model of its values, and for which you can make a mathematical model of how they combine, you can use a Monte Carlo analysis to come up with a statistical model of how the variables might work together. I don't know of any system that can make a general mathematical model of how the variables might combine, so a Monte Carlo analysis requires writing the core engine of a program for each problem.

Here's an absurdly simple example, the sort of thing that's commonly checked out with a spreadsheet.

Say you have a stack of 6 bricks in your assembly. All the bricks are nominally the same thickness, but there are three brick factories where you buy them, and of course, bricks aren't all =exactly= the same thickness, and the bricks from each factory tend to be a little different. You're going to make a million of the assemblies, so you want to know what you can expect the height of the tallest, shortest, and average stack will be. You'd also like to know what the odds are that the stack will be higher than some "magic height" where it won't fit.

You measure a bunch of bricks from each factory, and calculate the mean and standard deviation of each factory's output. If you're clever, you also make a histogram of thicknesses and see if the distribution matches (within reason) the "normal" distribution (Bell curve, Gaussian distribution).

Then you write a program that takes into account the number of bricks coming from each factory, and each factory's distribution, and roll the dice to make, on

paper (or computer, whatever), a large number of brick stacks. Say 20,000 just for laughs. You put the results into a histogram and report the mean, standard

deviation, min, max, number over "magic", and so forth.

The advantage of simulation is that you can tinker with the variables.

For Atlas 4, I wrote a Monte Carlo simulation of where the tracks would be on the disk. I used 34 independent variables, and the "assembly" used a lot of trigonometry to account for the angles as the actuator rotates in going across the disk, and for various "tilts" that happen. I used Borland Pascal 7 to do the job, with objects. The simulation ran at about 100 assemblies per second on a Intel Pentium-90, so you could simulate 50,000 assemblies in less than 10 minutes. It took a couple of hours to make sense of the results, of course.

We ran 14 different sets of input variables before we were happy with the answers. It took me a month to write the program (I was actually rewriting a similar one that I did for Viking 1), and about 3 weeks to go through the analysis loops.

The engineer who did a similar analysis for Viking II used an Microsoft Excel add-in, and he used to let the program run overnight on a Intel Pentium-200.

AD:

Can you discuss what is involved in the testing process for a hard drive? What basic tests are absolutely required to be passed before shipping the drive.

JT:

People make careers out of testing hard drives. There are the various engineering and qualification tests that each product has to pass before it's "shippable", then the detailed production tests that each drive has to pass before it's shipped.

Engineering and qualification tests are by no means identical, but I'll lump them for an SST-altitude view. I'm a mechanical guy, so I may miss some electrical or software testing in this list. It's not that I don't care, just that I'm ignorant of many details outside my specialty.

Operating and non-operating shock and vibration performance. Non-operating tests look for physical damage. Operating tests look for error rates and performance degradation in addition to physical damage.

Four-corner tests. Drive performance is measured at various combinations and rates of change of temperature and humidity, ranging generally about 5C beyond the specified temperatures, and usually some amount beyond the specified humidity (it's a lot harder to control humidity precisely).

Altitude tests. Drive performance is measured from 200 feet below sea level to at least 10,000 feet above sea level. Flying height and flying height variation is particularly scrutinized.

Voltage limits. Drives are typically specified to run at plus or minus 5% of specified voltage. Testing is commonly done to plus or minus at least 10%. There's normally a test to find out how far off you can go before the drive fails. All combinations of high, low and nominal 5V and 12V are tested.

RFI/EMI tests. The drive's electronic emissions are measured, and its susceptibility to external electromagnetic fields is measured.

Start/stop reliability. Samples are started and stopped massive numbers of times. Starting current, error rates and acoustics are measured at intervals. For a 40,000 start/stop spec, about 1000 drives are spun up and down maybe 80,000 to 100,000 times each. The test takes months.

Acoustics, both idle and seeking, both "new condition" and after various torture sessions. If a drive fails, it can be very hard to find out why and what to do about it. I've probably spent a total of 5 years working on acoustical issues, interspersed with other tasks.

Contamination measurements. This is usually done with drives that have been going through 4-corner or some kind of reliability testing. Test results can be incredibly baffling and hard to interpret and hard to figure out what to do.

All the various interface tests (data rates, error rates and so forth). There are many such tests, and I'm afraid I just don't know much about them in detail.

Latch reliability. All modern drives include some kind of a lock to keep the actuator parked in the landing zone while it's stopping and starting. There are several kinds, and many variations of each design. The lock has to keep the actuator parked while spinning down and not allow any combination of shocks and accelerations to let the heads move out of the landing zone while power is off. Both linear and rotational shocks and accelerations are tested. This is one of the most difficult tests to pass and one of the most hated assignments for a mechanical engineer.

TMR measurements and other servo performance measurements. This is a critical item. Servo performance is subject to strange failures on totally unpredictable combinations of seeking and external influences. Servo engineers have as hard a life as mechanical engineers! I've given you a very cursory discussion of TMR, and it shouldn't be hard to dream up dozens of tests from that, if you have a sufficiently evil mind. <g>

Thousands of drives are run for a few thousand hours each and power consumption, data-handling parameters, and reliability are measured. The final reliability test is so stringent that one hard data error in a couple of thousand drives, over a thousand or more hours each, can halt the program. Such a thing may happen once in two out of three development programs. There's hell to pay when it does!

Production testing.

A small percentage of new drives are destructively tested for non-operating shock, internal cleanliness, and such things.

Samples are measured for acoustic performance.

The rest of these tests are 100%. It's typical for such tests to take about 8 hours. The 36-GB Quantum Atlas 4 drive needs about 20 hours to do its

testing. That's partly because the error testing takes time directly proportional to the number of disks and partly because that drive gets unusually stringent testing because of its intended market.

Every head on every drive is measured for its reading and writing properties (amplitude, resolution, overwrite capability, PW50 [a measure of how cleanly a transition can be read], and nowadays some MR characteristics that I don't recall. The drive maintains tables of these parameters by head and zone on the disk. There are typically 16 data rate zones.

[Note by JT: Look at the article about "Disk Layout - Data Sectors and Servo Sectors" for more information about zones.]

[Note by AD: The data are kept in a "secret room" which John Treder will explain a bit more about:

A drive might have 4 disks, 8 heads, and 16 "data rate zones", each of which may have a different number of data sectors. Each head and each zone will have some of its critical read/write parameters measured during the factory test and stored away. That means a dozen or so tables of 128 values each (probably integers) to be kept somewhere. You don't want to put it in ROM, because it would be too expensive to have a unique ROM for each unit you build. And a hard drive is designed to hold variable information. So you simply keep all kinds of running and testing data on the drive. You also keep a good deal of the drive's operating code on disk, and page it to the drive's RAM as needed.

All that stuff is stored in "extra" tracks outside the user's data space. That's our "secret room". The extra tracks are formatted exactly the same as regular data space, it's merely on tracks -1 to -28 (or whatever), and they're part of the outermost data rate zone. There are usually 25 or 30 tracks "reserved" for that purpose, so a hard drive has room to store perhaps 20 or 30 MB of private programs and data.]

Every surface is scanned for media defects and hard error locations are re-allocated. There are algorithms used for "scratch-fill" to eliminate sectors between detected errors; those sectors are likely to have errors that just didn't quite get detected. Several hundred hard errors per surface is normal. The test includes deliberately moving the heads off the track center to find scratches or pits "between" the tracks.

Actuator latch opening and closing speed is tested.

Data rates and soft data error rates are measured by head and zone. Even if a drive may pass overall, it can fail on a detail.

Servo parameters such as raw seek times, settling times, stability parameters, and quality of the recording of the servo data are measured. Information about how much current it takes to stay on track, and what the relation is between seek current and acceleration (the torque constant) for several places across the disk is stored in the "secret room".

At Quantum, the production testing is done in an environment that's roughly equivalent to what a typical operating environment might be--ambient temperature around 40C, humidity whatever it is (factories are in Singapore, southern Japan and Ireland, so humidity is generally high), and about 100 drives running in a test cell, 10 drives per shelf, sort of bouncing around.

Criteria to pass the tests are always more stringent than the specs. That's been true everywhere I've worked. I thought about the possibility of trying

to give specific numbers for these tests, but as I thought about it I realized that the test criteria change so fast that what I know is certainly already obsolete.

AD:

Has there been any consideration to make flash ROM (similar to what most modems have) a standard for hard drives? Or is that considered too dangerous?

JT:

During development and often leaking into the early production drives, there's a flashable ROM on the drives. It's replaced as early as possible with hard-coded ROM to save costs. And afterwards, if the ROM needs changing, it's an earthquake-level task.

So it isn't danger, it's a combination of \$\$\$ and tradition.

AD:

If cost was not an issue what kind of hard drive would you design for your own use?

JT:

10K RPM, 2 1/2" disks, about 20 GB, with two complete drives, striped, in one housing. It should be able to put about 60 MB/s across the interface, continuously.

It won't happen. It's WAY too expensive, but technically pretty easy to do.

AD:

Any thoughts on other storage technologies such as CD-R/CD-RW, Magneto-Optical or DVD? Do you think these technologies will replace hard drives as the primary storage medium?

JT:

I don't think CD and its derivatives will replace hard drives because the physics in the way they write, especially, is slower than hard drive magnetics. MO has no advantage over hard drives in speed or data density.

However, sometime before 2010, hard drives will hit the wall in terms of data density, and at this time I don't know of a way around it. That's the first time I've had to say that in my hard drive career. In the past, I've been able to perceive one or more ways around some supposed barrier to speed or capacity. When the data density barrier (technically it's called the paramagnetic limit) is reached, the only candidate I see for replacement today is some development of flash ROM. They need to cut cost by an order of magnitude and improve writing speed by a couple of orders of magnitude. Those are formidable challenges!

AD:

Any common misconceptions about hard drives that end users have that you would like to clear up?

JT:

The biggest one is that you'll wear out the bearings by letting the drive run. If you leave a drive running continuously, there's roughly a 1% chance that a bearing will fail in the first 7 years.

The only other one is the argument about whether to leave it running or shut it off. I just said it won't hurt to leave it running. Well, the standard test for starting and stopping ends up with a 0.3% chance of a drive failing to start in the first 20,000 starts.

So my advice is, leave it running if you like, shut it off if you like. It doesn't matter. If your drive fails, it isn't because of your choice in that matter.

AD:

Did you work a 40 hour work week during the design process?

JT:

As I said before, Ha,ha,ha,ha!! During the heart of a project, 60 or 70 hours was pretty much normal. When there's a crisis, or during a build, 80 to 100 hours of actual working (not just being there) is what you do.

It's funny, of course. One guy will be busting butt, and the guy in the next cube will have nothing out of the ordinary on the fire. Yet engineering is essentially an intellectual sport, so you can't just hand off half your work when you have a crisis.

AD:

How much documentation was produced for the Atlas 4? An estimate would be fine.

JT:

Depends on what "documentation" means--but let me see--print docs, maybe a pile 20 feet high, if you don't count all the drafts.

The hard-copy documents I kept in my file filled an entire file drawer for each of three programs I worked on at Quantum. I threw away a lot more paper than I kept.

There were 200+ mechanical drawings, with 2 to 15 revisions each (average maybe around 5 revisions). I didn't have a complete solid model file of the drive, my solid model directories for Katana (Quantum's internal project name for Quantum Atlas 4 (SCSI) and Fireball +KA (IDE); they're the same except for the interface) never ran more than about 200 MB. Two other engineers kept solid model directories, too.

The firmware manager had a graph on his office wall about size of the code files--I think it peaked around 4 or 5 gigs.

E-mail & phone-mail messages, I have no idea. My E-mail was constantly overflowing my 4 MB allocation--about 500 messages or so. I had to purge it once every couple of months. Managers had more space <g>.

In other words, lots and lots of docs.

AD:

What is the estimated cost to bring a new hard disk to market from start to production?

JT:

These days, in the range of \$50 million.

AD:

Can you discuss a bit about the ATM (Automatic Teller Machine) deposit mechanism that you designed (20+ years ago)?

JT:

20 years is a long time! The problem was to accept deposit envelopes of various sizes and shapes and thicknesses, maybe containing coins, print an identification number on them, pass the UL test for theft resistance, work when it's pouring with heavy rain, fit in the available space, be easy to maintain, and cheap. In general, the usual engineering challenges. <g>

The hardest task was to pass the UL break-in test. The tester was a massive, muscular fellow armed with punches, crowbars, sledgehammers, long grabbers, and other tools. He could study the deposit system for as long as he wanted to, inside and out, before he began his attack. He had half an hour of actual "breaking in" time to try to fish out an envelope. That half hour didn't have to be contiguous. He could bang on the depository, then go around and see if his attack was working. Another UL person timed him with a stopwatch. We passed, barely.

It was also difficult to come up with a reliable envelope printer. I eventually designed a sort of rotary rubber stamp that printed the number every couple of inches along the envelope. If there were coins, it seemed from our testing that there was always a way to make out the number, maybe combining a couple of partial printings.

AD:

Since you have done race car driving via Sports Car Club of America (SCCA) road racing do you ever play with car racing simulations/video games?

JT:

I've tried a couple, also tried a couple of coin-op games. They're boring.

One of the arcade games had pretty good visuals, comparable to the in-car cameras you see on TV occasionally. The problem I have with all those things is that you don't get the physical feedback. 1.5G+ cornering forces, 1G+ braking, etc. If you get a flat-spotted tire it can literally cause you to see double. The simulations don't do that stuff. Also, they're generally over way too quick. If you imagine more intensity than an arcade racing game, then have it last for 45 minutes or so, that's what you do. I'm a skinny sort of guy, 5' 8" and 140 lb, and I used to sweat off 5 lb or so in a 40 minute session. It was fun (the most fun you can have while dressed), it was intense, and it required total commitment, not only on the race track, but in preparation too. When I was no longer prepared to give the commitment, I retired.

AD: Thank You.

Mr. John Treder has also written more details on hard drive history, design, and performance issues which are included in the following sections.

=====

History of Hard Drives

=====

The hard disk drive was invented by some IBM engineers working under Rey Johnson at IBM in San Jose, CA, in about 1952 to 1954. I worked at IBM from '65 to '81 and got to meet and work with some of those men - Rey Johnson, John Lynott, Don Cronquist, Bob Schneider and Lou Stevens come to mind right now.

In 1965 (I think that was the year) a number of engineers left IBM (they were known as the "dirty dozen" within IBM) and founded Memorex. Al Shugart, one of them, later left Memorex and founded first Shugart Associates where the 5 1/4" floppy disk drive was a major product, then Seagate Technology, which effectively started today's industry of small hard disk drives.

The early drives almost all had linear actuators, that is, they moved the heads across the disks in a straight line, using a carriage with wheels. It was only later that rotary actuators, where the heads are held at the tips of a comb-like array and they swing back and forth like a gate, became popular. Because the rotary actuator is cheaper, it's now the standard for all hard disk drives, and that's what I'll be talking about.

The first IBM RAMAC disk drive had a couple of dozen disks, each about 2 feet in diameter, and ONE head! The head was moved from disk to disk and back and forth on each disk using a system of cables and pulleys and stepping motors. The added speed of having at least one head for each disk surface, and of using both surfaces of each disk, soon became obvious, and drives began to look pretty "modern" by 1960, although they were vastly larger and more expensive. Whether the heads are moved in a straight line or swung in an arc, something has to provide the force and the control to move them and keep them in the right place.

Stepping motors, hydraulic actuators and voice coil motors have been used to provide the motive force. Stepping motors have a built-in capability to hold in one position. Hydraulic actuators and voice coil motors (VCMs) provide force, but can't hold a position with great accuracy. A rack with detent pawls has been used, but nowadays a servo system is used, with the positioning information recorded on the disks.

So you can have a disk drive with a stepping motor and you don't need a "servo" or you can have a disk drive with "servo data" recorded on the disks. A stepping motor is simpler (at least in concept) and cheaper, but it's slower in seek time because it isn't really very powerful, and it isn't capable of really, really fine precision. A VCM can provide enormous forces, but it needs control, which the servo system provides. However servo feedback systems are complicated and you have to pre-record the positioning data somehow (usually on one or more of the disks), and that takes up space that could be used for "real" data. On the other hand, servo systems can provide incredibly precise positioning.

The style of hard disk drive we use today began to emerge in the early '80s. I think it was Maxtor, under Frank Gibeau, where the first high-volume 5 1/4" disk drives with a rotary actuator, a VCM and a servo system were

produced.

In 1986, Finis Conner left Seagate and founded Conner Peripherals along with John Squires, and they built the first high-volume 3 1/2" disk drives. The first one, 40 MB, was called the "Fat 40". Not only did they popularize the new smaller "form factor", but they were the first to have an "embedded servo" or "sector servo" in volume.

Meanwhile, Quantum Corporation had been building 8" and 5 1/4" disk drives since 1980, and in the mid-80's they (actually I think it was Joel Harrison and Bill Moon) saw an opportunity with the 3 1/2" form factor and invented the "hard card", a disk drive on an expansion card that you could just plug into your AT. And that's how the IDE interface got started.

By the way, Quantum used a rather odd variant of the servo system for many years, where the servo information was actually very fine lines etched on a piece of glass attached to the actuator, and read with a photocell. It's actually more complicated, but that's a subject for another discussion.

Around 1990, laptops began to appear, and with them came the 2 1/2" form factor. I never worked closely with 2 1/2" drives, and I'm not very well versed in their historical development.

Way back in the old days, when the world was young and all, <g>, a single disk surface was reserved for "servo data" on disk drives that had a voice-coil actuator. Drives with stepper motors didn't need that stuff, and could justly claim that they didn't have to "waste" disk surface on "servo data".

Servo data, by the way, is information that's pre-recorded on the disk and specially formatted to make it possible for the drive to know where its heads are positioned. A stepper motor has, if no errors have occurred, the equivalent information built into its mechanical structure.

But during the late '70s and early '80s, techniques were developed that allowed the servo data to be written on the same surfaces that hold the regular user data. There were several schemes proposed and actually implemented, but the one that has taken hold is called "sector servo", where some number of regions on each data track are specially reserved for servo information. Because the sectors are physically coherent on each disk surface, they're commonly called "servo spokes". I've seen drives with as few as 64 spokes per revolution and as many as 128.

I believe the first major manufacturer to use sector servo was Conner -- sector servo and the 3 1/2" form factor were their keys to success when they started in (I think) 1986. [At the time, I was one of the majority of old disk drive people who thought they were headed for disaster. Later, I worked there. Shows how smart I am!]

While there may still be drives manufactured with a dedicated servo surface, I think the last major manufacturer to use one was Seagate, up until a couple of years ago. The first Barracuda drive had a dedicated servo surface. Later Barracudas, and I think all current Seagates, use sector servo technology.

DRIVE PERFORMANCE - TMR

One of the ways that disk drive engineers study performance is to measure TMR-Track MisRegistration. It's really quite a deep subject and servo engineers know more about it than I do, but I think I can give you an

"executive" overview and perhaps add a little bit to your understanding of how disk drives work.

First, I'll quickly review the basics of track layout. Data tracks are laid down in concentric circles in a hard disk drive, not in a spiral as on an LP record or a CD. In a modern "embedded servo" drive, there are anywhere from about 70 to about 120 "servo bursts" per revolution, on each track. These servo bursts contain the special data that the head reads and the servo system firmware interprets to determine which track the head is closest to, and how far off "track center" the head's center is. For the space between servo bursts, the system's inertia and the current supplied to the actuator by the servo system keep the head in position, or "close enough", and that's where TMR comes in.

TMR-Track MisRegistration is just a fancy phrase for "error". It refers to where the head's read gap or write gap is, relative to where you want it to be. For convenience, I'll just say "head" from now on, although I'll actually be talking about the spots on the head where reading and writing happen. And in this essay, I'll be talking about "static" TMR, the errors that occur after the head has been "following" a given track for long enough that errors resulting from seeking from one track to another have faded out. I won't be talking about the causes, I'll just be talking about the errors themselves and how engineers think about them.

There are several flavors of TMR. First, there are "static repeatable runout" and "static non-repeatable runout", usually referred to as RRO and NRRO. Those refer to the deviation of the head from the theoretical perfect circle of the track. In this case, static doesn't mean the disk isn't spinning, it means the head isn't seeking. RRO is "phase locked", that is, the head is off track by the same amount at the same point on the disk, each revolution. So we can talk about "once around" or "5 times per rev" RRO, which for a drive spinning at 7200 RPM (120 Hz) would occur at 120 Hz or 600 Hz. NRRO has characteristic frequencies, but they aren't locked to a particular location on the disk. We speak of "506 Hz" and "570 Hz" NRRO frequencies. (Those happen to be two frequencies of particular interest to a certain Viking engineer.)

Then there are several factors that as a mechanical engineer I just lump together as "servo TMR". The servo system has white noise in it, several resonant frequencies, and various filters, including several "notch filters" to hide the effect of mechanical problems. The notches add certain problems of their own.

Those TMR factors are important to an engineer, but they don't have anything directly to do with your data. There are three main TMR factors that are measured using data. They are "Write TMR", "Read TMR" and "Write to Read TMR". (Write to Read TMR is often called just TMR, because engineers tend to give long, clumsy but accurate names to things, then use the same handy short name for several different phenomena, just to confuse themselves and others.)

Let's think about what happens when data is written and read back. As a head passes a servo burst, the "tracking error" is measured and a correction factor is calculated and the current to the voice coil is adjusted a little bit, if necessary. Then we write a sector that's between this burst and the next. The data is written where the head is, generally not on the exact center of the track. The distance between the theoretical track center and the center of the actual data is called "Write TMR". Some time later, the head tries to read that data. The head passes the servo burst, the tracking error is calculated, and the head tries to read the sector that

was written. But the head is generally not only not on the center of the track, it's generally off by a different amount than it was when the data was written. The distance between the theoretical track center and the center of the read element is called "Read TMR". The total error, the distance between where the data was written and where the head is when the data is read, is called "Write to Read TMR". Write to Read TMR is very important, because it represents the sum of most of the things that can cause data-handling problems inside a disk drive.

If W/R TMR is so important, how can it be measured? After all, it's defined in terms of data, not servo information. Well, it happens that if you write known data patterns and calibrate the results, you can measure parameters that allow you to calculate W/R TMR. What is written is low-frequency information and high-frequency information. The input data isn't "all 0's" and "all 1's" because the PRML coding won't translate that into constant frequency, but the idea is to get the effect of 0's and 1's. If you have constant-frequency information, it's easy (or so I'm told) to measure its amplitude accurately, and disk drive read channels do measure the amplitude (for setting AGC). It's also possible to command the servo system to move the heads X% off the track center (where X% is a percentage of the track-to-track distance). By reading the data amplitude at various amounts of deliberate off-track, you can calibrate the amplitude against distance. And you can do that even when all the errors that you're trying to measure are happening, if you just take enough measurements under carefully controlled conditions and find the average.

OK, so that's what TMR is. Now, what does it mean to the user? If total W/R TMR is too large, you may not be able to read the data that you wrote. If Write TMR is too large, you may erase the data that was written on the next track. You aren't likely to contemplate either of those eventualities with perfect equanimity.

The servo system, as I said, is measuring the head's deviation from the track centerline at each servo burst and continuously correcting the errors that it finds. If the head drifts too far off center for it to be safe to write, the servo system posts a "write inhibit", telling the system not to write data until further notice. If the head drifts even further off center, it posts a "read inhibit" telling the system that if it reads, it'll likely not be good data. As the head returns toward track center, if its position and velocity are both within spec limits for a spec length of time (# of servo bursts), "read inhibit" and, using tighter specs, "write inhibit" are removed. Writing is controlled by more stringent specs than reading because if you write over some other existing data, there is no recovery at this low level. The data is just gone and you don't even know it! (Again, at this very low level.)

So what are the spec limits? I won't talk about velocity and time limits, those both vary greatly between different companies and different products. Distance off track is usually spoken of as a percentage of the distance between track centers. Write inhibit is usually set if the head gets more than about 10% off center, read inhibit is set at 15% or so. The new MR heads have separate read and write elements, so the old tape mantra of "write wide, read narrow" has been implemented, and there's a tendency these days to open up the read limits a bit.

For the newest drives, with track densities in the range of 10,000 TPI, that implies that the servo system keeps the heads on track within plus or minus 10 microinches for writing. But those numbers are for safety. For performance, you have to be better. If Write to Read TMR is less than 5% of track 99.7% of the time, you'll have a top performing drive. Yes, that

implies staying within 2.5 microinches of track center!

=====

SEEK TIME

=====

I offered to post a "short course" on hard drive seek times, and here it is. There's some historical information that may help some of you to understand what I'll be talking about.

I'll get a little bit technical, but there isn't any math more complicated than elementary algebra. I'll state various versions of Newton's laws of motion and Faraday's laws of electromagnetics without proof.

The path I'll follow is essentially how disk drive actuators were designed about 15 years ago, and it's still how the first cuts are done before optimization. I'll discuss a 3 1/2" disk drive because it's the most popular size today, and I'm working on one now so all the numbers are easy for me to remember.

Definitions

For reasons having to do with laziness and tradition more than anything else, a strange mixture of English, CGS and SI units are used in actuator and VCM design. I'll give some of the more common values and definitions here, so I won't have to interrupt the flow.

Actuator:

The assembly inside a disk drive that holds the heads and moves them back and forth over the disks.

VCM:

Voice Coil Motor. The assembly of magnets, steel and coil of wire that provides the forces that move the actuator back and forth. The coil is mounted on the actuator, and the magnets and steel "pole pieces" or "mag plates" are mounted on the drive's base casting.

Head:

Specifically, the small piece of ceramic that holds the read/write elements and the air bearing surfaces; that is, the stuff that makes it possible to read and write the magnetic information on the disk. Generally, the assembly of that ceramic piece with its stainless steel suspension mechanism. To a disk drive engineer, a "head" is always more specific than an "actuator" and a "slider" is more specific than a "head", and "ABS" (air bearing surface) and "R/W (read and write) elements" are more specific still.

Pivot:

The ball bearing assembly on which the actuator pivots (swings like a gate). The heads are on one side of the pivot, and the VCM is on the other side. (That isn't theoretically required, but it's almost always the case.)

OD:

Outside diameter (Sometimes I'll talk about the OD radius; I hope you'll forgive me.)

ID:

Inside diameter

Stroke:

The distance across the disk that the heads can move. "Full stroke" is from the OD to the ID or vice versa.

1 inch = 1 in = 1" = 25.4 mm

1 mil = 0.001 in = 25 microns (roughly)

1 microinch = 1 uin = 0.000 001 in

1 micrometer = 1 micron = 0.001 mm

1 micron = 40 uin (roughly; the exact number is 1 micron = 39.37 uin)

1 Tesla = 10000 gauss

1 ips = 1 inch per second (traditional unit)

1 g = 32.2 feet per second per second = 386.4 ipss = 9800 mm/sec² = the acceleration of gravity at the earth's surface

1 newton (unit of force) = the force required to accelerate a mass of 1 kg at 1 m/sec² = 0.448 lb, roughly. Acceleration is the rate of change of velocity, so it's distance divided by time, divided by time again.

a' = angular acceleration. The units are angle per second per second, usually radians/sec² in engineering.

2*pi radians = 360 degrees = the angle for one complete circle. (1 radian = 57 deg, roughly)

Seek Profiles

What we're going to do is move the heads across the disks. It seems that most people are interested in doing that as fast as possible, so a good place to start thinking about actuator and VCM design is how you want it to move.

A 3 1/2" disk has an OD of 95 mm and an ID of 25 mm. However, you have to hold the disks to the spindle, and the clamp and spacers often have an OD of 31 mm, so theoretically, you have a maximum usable space from 15.5 mm radius at the ID out to 47.5 mm radius at the OD. You need to have a place to park the heads when the drive is stopped, and that is always at the ID, because that's where the motor will need the least torque to start up. It turns out that after monkeying around with tolerances and so forth, the outer data track will be at a radius of about 46 mm and the inner data track at a radius of about 21 mm, so the "data zone" is pretty close to 1" wide.

Some IBM engineers published a theoretical analysis in the early '60s that showed how a 1/3 stroke seek, that is, a seek across 1/3 of the data zone, takes an amount of time equal to the average time for a large number of random-length seeks. Whether that conclusion is valid depends, of course, on whether their assumptions are valid. Well, they made very good assumptions.

This

short course uses the same assumptions they did, and we'll calculate 1/3 stroke seek times and have full confidence that they are equivalent to random access seek times. Real disk drives today don't exactly follow the assumptions, and if you measure a modern drive, you'll probably get a slightly different answer, by 0.1 msec or so.

If the data zone is 1" wide and the average access time is defined by 1/3 of that, then we only have to study the access time for a 0.333" stroke. I'll round that to 0.3" for discussion.

How fast can you go for 0.3"? I'll start by stating an "effective speed limit" of 250 ips. It isn't smoke, you'll see a number pretty close to this later. Let's calculate a "Roadrunner" seek. Suppose acceleration were meaningless, and you could be going as fast as possible, instantly, and then you could stop instantly. Boinnggg! Beep-beep!!

Newton said (among other things) $t = s/v$. If $s = 0.3$ " and $v = 250$ ips, then $t = 1.2$ msec.

Hmmm. Pretty fast. But what does it mean? If a Top Fuel dragster can be going 300 MPH at the end of the quarter, then he ought to be able to do the quarter in 3 seconds, using the same $t = s/v$ equation. Well, they can't do that (but they're amazingly close) and we can't do 1.2 msec seeks either. And it's easier for a dragster than for a disk drive; in a disk drive you have to both start and stop, while a dragster isn't timed for stopping. A disk drive isn't Roadrunner. It can't be going at maximum velocity instantly. It has to accelerate and decelerate. Let's calculate a "dragster" seek. We'll accelerate like a dragster and then decelerate at exactly the same rate. This is called a "bang-bang triangular seek" to we few cognoscenti. Since it's symmetrical, all we have to do is calculate the first half, then double the time. The equation according to Sir Isaac is $t = \sqrt{2s/a}$. We have to pick a distance (0.15") and an acceleration. I'll grab 50 g's, not quite out of the air. Multiply 50×386.4 to convert acceleration from g's to ips and make everything come out right. The time (2t, remember) is 7.9 msec!

Does this mean that we're flinging the heads around at 50 g's or more in a cheap \$400 4-GB disk drive? Yes. Actually, it's worse.

Stay tuned for the next thrilling episode!

Previously I introduced the idea that you need to accelerate and decelerate the heads when you seek, but I didn't show any particular reason for choosing the velocities and accelerations that I used in the examples. Today, I'll let you peek over my shoulder as I get the specifications and do the first rough design for a new (hypothetical) disk drive.

The boss called a meeting and said, "The bozos in Marketing just said our new drive has to have an average seek time of 8 msec. I don't want to be caught short when they change the specs, so we'll design for 7 msec. OK?" At this point I didn't have any data to tell the man to take a long walk off a short dock, so I merely looked very grumpy. Then I went and talked to the servo engineers, and they said they'd like to have 3 msec of settle time. I said, "Holy cow! Do you really want me to do a bang-bang of 4 msec?" They said, "Sure!", because it wasn't their problem.

So I went back to the cube and fired up the rusty old spreadsheet.

1/3 stroke distance $s = 0.3$ "

time $t = 4$ msec

The acceleration $a = 2s/t^2$ where s & t are half the time and distance, because I'm calculating the acceleration half of the seek and assuming that the deceleration will be just the same.

So using $s = 0.15$ and $t = 0.002$, $a = 75,000$ in/sec² = 194 g's!

The maximum velocity, when you stop accelerating and start decelerating, will be $v = at = 150$ ips = 8.5 mph.

Now, the servo guys don't care about my problems with acceleration, but they do worry about maximum seek velocity. So I went back to their main man and said, "How does 150 ips sound?" He said, "Uckk! I'll get back to you." I

expected some such response, and I have my own less detailed version of their calculations tucked away in my spreadsheet.

Here's my version of the servo problem:

The new 7200 RPM drive is going to have about 10,000 tracks per inch. Embedded servo systems have roughly somewhere between 75 and 100 sectors per revolution. We call them "spokes" because they look like slightly curved spokes on a wheel when you "develop" a disk and make the magnetic patterns visible. I'll assume that this drive will have 90 spokes. So, if the drive is sitting on track, there will be $120 \times 90 = 10,800$ servo bursts per second, or a servo burst will come along every 92.59 microseconds.

If I'm seeking at 150 ips, I'll be getting $10,000 \times 150 = 1.5$ million track crossings per second, or doing a track crossing every 667 nanoseconds, that is, every $2/3$ of a microsecond. When I'm seeking at 150 ips, I have $1,500,000 / 10,800 = 138.9$ track crossings between servo bursts. Now, that's an awful lot of tracks to cross by "dead reckoning", so I expected the servo guy to propose a speed limit. Sure enough, I got a call saying, "How about 120 ips?" I figured he'd probably try to add a few more spokes and have a maximum of about 100 "blind crossings" at full speed.

Well, back to the spreadsheet. Now I'll accelerate to 120 ips, "coast" at speed for a while, then decelerate. The distance $s = 1/2 at_1^2 + vt_2 + 1/2 at_1^2$, the velocity $v = at_1$ and total seek time $t = t_1 + t_2 + t_1 = 2t_1 + t_2$. If I substitute the velocity into the distance I get $s = 1/2 vt_1 + vt_2 + 1/2 vt_1 = vt_1 + vt_2$. Rearranged, $s/v = t_1 + t_2$. Subtracting, $t_1 = t - s/v = 1.5$ msec, so $t_2 = 1$ msec and $a = v/t_1 = 80,000$ ipss = 207 g's.

Now I can tell the boss:

"Here's the story to get a 7 msec seek time with the servo requirements of 120 ips max velocity and 3 msec settle time. We accelerate at 207 g's for 1.5 msec, coast at 120 ips for 1 msec and decelerate at 207 g's for another 1.5 msec. Then we settle for 3 msec. I haven't done the VCM numbers yet, but I'm afraid the current to get that acceleration will be pretty high. While I'm doing that, can you maybe campaign for a little relief on the settle time?"

Next Problem: How to design a VCM without knowing all the details, and maybe, some negotiation on the specs. Previously I came to the conclusion that to build a disk drive with 7 msec average seek time that includes 3 msec settling, I need to provide a VCM that will generate 207 g's of acceleration of the heads. I'll abandon the story format for now and engage in a short lecture on Newtonian physics and how Faraday's laws of electromagnetics interact.

Up to now, I've been describing the seek parameters as if the head were moving in a straight line, but in a rotary actuator the heads actually move in the arc of a circle. Fortunately, all we have to do is wrap the straight line around the circle.

Newton's second law of motion is usually described as $F = ma$ for linear motion, but it's equally valid for rotation about a pivot. The rotary version is $T = Ia'$, where T is torque, I is the polar moment of inertia about the pivot and a' (a prime, because I can't be sure everybody has an alpha symbol in their font) is the angular acceleration.

Torque is force * the distance from the pivot. $T = rF$.

Polar moment of inertia is, officially, the second moment of mass about the axis. Using integral calculus, it's $I = (\text{integral}) r^2 dm$. That means that each particle of mass in the object contributes to inertia as the square of its distance from the pivot.

a' , angular acceleration, is measured as angle per second per second instead of distance per second per second. The conversion is $A = ra'$.

Hey! All I need to know to convert from linear to angular is the radius r ! Oops. To get the required torque, I need to know the moment of inertia. The way real engineers fill in a number that isn't known yet for a new design is to use a number that's "close" from an existing design. The radius r is going to be 52 mm, because that's pretty close to the numbers for many current disk drives. The inertia is going to be $4400 \text{ gm}\cdot\text{mm}^2$ because that's realistic, too.

The angular acceleration $a' = A/r = 39069 \text{ radians/sec}^2$.
 The torque is $Ia' = 39069 \cdot 4400/9800 = 17541 \text{ gm mm}$. The 9800 is the acceleration of gravity, used to convert grams mass to grams force.
 I'll also need the number in Newton meters:
 $T = 4400 \cdot 39069 \cdot 1\text{E-}9 = 0.172 \text{ N m}$.

Maybe the torque number will mean more to some of you if I express it as 1.5 lb in, that is, a force of 1.5 pounds at a distance of 1" from the pivot. The heads are at a radius of about 2", so we're putting a sideways force on the heads of 3/4 pound to accelerate them at 200+ g's. Yes, they're very light!

In a disk drive, the required torque is provided by a VCM (Voice Coil Motor). It's a variation of the same device that makes loudspeakers work, hence the name. The principle of operation was noticed late in the 19th century by Michael Faraday. If you have a length of wire in a magnetic field and you run a current through the wire, you get a force. In physics, the relation is expressed as a vector cross product, but in disk drive design we go to great pains to keep everything at right angles, so we can just use the algebraic equation $F = B L i$. F is the force, B is the magnetic field strength, L is the length of wire within the field, and i is the current. The flux B is applied across the wire, and the force appears across the wire, too, at right angles to the flux. Here's an ASCII picture (I hope it helps):

```
===== --> direction of current in wire
^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^
flux going upward
```


The force is going into the page (right hand rule).

I want to calculate the current. Rearranging, $i = F / BL$. The equation requires force, and I know torque. To convert, we divide the torque by the radius to the "effective point of application" of the force. So, $i = T/BLr$. To make it work right B has to be Tesla and T has to be Newton meters. I'll use a radius of 20 mm (about 3/4") from the pivot to the center of the coil. If $B = 0.9$ Tesla, $L = 20$ mm, $r = 20$ mm and T (above) is 0.172 Nm, then $i = 478$ A!

Holey Moley!

Actually I just did that to impress everyone. The calculation above assumes one length of wire 20 mm long in the magnetic field. Fortunately, we can wind a coil with lots of turns and use less current. Each turn contributes

its own piece of force. In addition, we use 2 identical lengths of wire for each turn, passing through the magnetic field. It's time for another ASCII picture!



I'm not clever enough to draw the magnets in their actual pie shape, so they're shown square, and the coil is shown as a trapezoid. Also, there's a piece of steel behind everything, then a magnet below the coil, then another magnet above the coil, then another piece of steel. Imagine that the pivot is about where the angled legs of the coil point to. Current flows (say) clockwise around the coil. So with the current going up in the left leg and the flux coming out of the North pole of the left magnet, the force is to the left; on the right, the current is going down but the flux is going into the South pole of the right magnet, so the force is still to the left.

Real VCM's may have from about 130 to maybe 200 turns of wire in them. Assuming that our (not yet designed) VCM will have around 150 turns, we'll need to get $478 / (2 \cdot 150) = 1.59A$ of seek current.

OK, now I can go back to the boss and tell him, "We're OK for your ridiculous 4 msec bang-bang time if you'll let me have at least 1.6A seek current." Fortunately, I get along well with my boss and he can stand a bit of ragging!

So there's a first rough pass at the design of a disk drive actuator for a specific seek time. There will be several passes around design calculations similar to the ones I've outlined, adjusting maximum current draw, maximum velocity, settle time allowance, and even the seek time spec, until a design emerges that's "equally unsatisfactory to everybody".

Those of you who know about this stuff will notice that I blithely ignored a couple of very significant factors--back EMF and inductance. Next, I'll try to discuss how they fit in (mostly, they just mess up the beautiful symmetry of these simple calculations) and I'll also try to get into settle time (from a distinctly high-altitude perspective).

Previously I finished calculating the raw bang-bang seek time of an actuator and figuring out how much current the VCM was likely to draw.

To review, the specs came out:

Seek time = 7 msec	(given)
Settle time = 3 msec	(given)
Raw seek time (bang-bang) = 4 msec	
Accelerate time = decelerate time = 1.5 msec	
Coast time = 1 msec	
Radius, pivot to head = 52 mm	(given)
Radius, pivot to center of VCM = 20 mm	(given)
1/3 stroke seek distance = 0.3"	(given)
Tangential acceleration at head = 207 g's	
Maximum tangential velocity of head = 120 ips	(given, after negotiation)

Angular acceleration of actuator = 39069 rad/sec^2
Inertia = 4400 gm mm^2 (given)
Torque for acceleration/deceleration = 17541 gm mm
Magnetic flux in gap = 9000 gauss (given)
Number of turns in coil = 150 (given)
Seek current = 1.59 A

It looks like a pretty complete lineup, even though I just tossed in some of the dimensions and VCM specs without any real explanation. If you tear apart a drive and play with a ruler and a gauss meter, you can verify that the numbers are "in the ballpark", or you can trust me!

I mentioned before that there are a couple of flies lurking in this fragrant ointment! The first and most important is back EMF. When a current is passed through a wire that's in a magnetic field, a force appears on that wire, and that's the $F = Bli$ equation I used yesterday. It's also been noticed that when a wire is moved through a magnetic field, a voltage appears across its ends. The equation is $E = BLv$, where E is voltage and v is velocity. The voltage is called "back EMF" because the voltage appears backwards to an engineer who just used current to generate a force and make the wire move; the voltage opposes the flow of current.

How does this affect the VCM? Well, at 120 ips at the head, the VCM coil is moving at an average of $120 \times 20 / 52 = 46.2$ ips. Converting everything to SI units and remembering that we have 2 passes and 150 turns for L , we get

$E = 0.9 \times 2 \times 150 \times 0.02 \times 46.2 / 39.37 = 6.33 \text{ V}$. By the time I get to full speed, I've lost over 6 volts of my driver! In a 3 1/2" drive we run the motor and VCM with 12V. After taking power supply tolerance and driver FET drop and resistance of the leads into account, we have about 9.5V available at the coil, and its wire size will be chosen to provide a resistance that just about allows the maximum start current. In this case, that's about 5.9 ohms (probably a little less in practice). Now at top speed I only have $9.5 - 6.33 = 3.17 \text{ V}$ available to push current, so I only get $3.17 / 5.9 = 537 \text{ mA}$, max. What happens is, as the actuator accelerates, it develops back EMF, so it gets less current and doesn't accelerate as fast any more. In fact, the acceleration curve is a decreasing exponential.

On the other hand, what happens when you switch the current around the other way to decelerate? All of a sudden, the back EMF is working in the SAME direction as the driver, and you have MORE voltage to push current through the coil! In our example, you get $(9.5 + 6.33) / 5.9 = 2.68 \text{ A}$ of current. Of course, as you slow down, the current decreases until you get back to the 1.59A static value.

In any case, the power driver chip guy wails, "Oh, my poor smoking FETs!" (They do tend to get dramatic at times.) So in real life, you must accelerate slower than the simple calculation, but you can decelerate faster. In theory, it all comes out even. In practice, the driver doesn't mind being saturated for acceleration (if the coil resistance is appropriate), but it usually needs to "control" on deceleration for reasons having to do with the servo code, as well as preventing the smoke from leaking out of the chip. So back EMF slows things down.

The other thing that messes up the tidy little calculations is inductance. When you pass a current through a coil of wire, you observe a property called inductance. The effect is that you apply voltage, but it takes a while for the current to rise to the value that you expected from measuring a straight piece of wire, and then when you remove the voltage, the current persists for a while. You can refer to the inductance as a quantity measured in a unit called henries (from Joseph Henry who discovered some of

this kind of stuff) or you can refer to a "time constant", the time for the current to rise to 63% of its final value. Henries (or millihenries) don't help me much, so I usually remember the time constant, and typical coils for disk drives have a time constant of about 100 microseconds or a little more. For those who'd like to calculate or measure, the coils in disk drive actuators are air core coils with an effective diameter of about 1". They're actually shaped like a piece of pie with the tip eaten off. 100 usec doesn't sound like much, but it's 6.7% of our accelerate time.

Settling Time:

One of the reasons that the initial specs for a disk drive have a large allocation for "settling time" is that when we calculate the raw seek time, then add the fudge factors for back EMF and inductance, the actual seek time eats into the settling allowance. You can figure that the actual seek time will be about 5% to 8% more than you figured. In a 7 or 8 msec drive, it's about 0.5 msec.

But settling is a real factor, not just a fudge. When you've just been decelerating at 200 g's for 1.5 msec, you've excited every vibration mode that the actuator knows about. The whole thing wants to ring like a bell. And despite the attempts of us mechanical men to make everything stiff, it's in the nature of metals to ring!

I suggested that my theoretical drive would have 10,000 TPI. Viking was rather conservatively designed; it only has 6432 TPI. The new IBM 2 1/2" drives have to have about 12,000 TPI if the reports in the papers make any sense (I haven't seen actual specs). Anyway, 10,000 TPI will be a commonplace within the next 6 months, and that translates to 100 uin from track to track.

A head has to be within about 15% of the track centerline to read data reliably, and if you allow a head to drift more than 10% off the track centerline while writing, you're taking very serious risks of destroying data on the next track over. Remember, the next track could also be 10% off closer to your track, and we want to write as wide as possible to improve the integrity of the data. We can't have very narrow tracks with large safety margins between them and get these huge data capacities; we have narrow tracks with even narrower margins! 10% of 100 uin is only 10 uin. It isn't hard to get a head to vibrating back and forth with a zero to peak amplitude of 10 uin!

To get the best possible seek time, the servo system is slightly underdamped. That is, as it comes up to the destination track, it's expected to overshoot slightly, then correct and wobble back and forth a couple of times. Ideally, the head should stay within the 15% on the first wobble, but it works out to be OK if it's on the second one. If you tried to land precisely on track, you'd have to approach more slowly, and it turns out that the slow approach takes more net time than the chance of an extra wobble.

Here's another case where I really wish I had a virtual whiteboard! This stuff would be a lot easier to see if I could draw you a picture, but my ASCII art abilities are sadly limited!

Before I mentioned that a servo burst might occur every 92.59 usec. Say it's 100 usec to make it easier to figure. That makes the sampling frequency 10 kHz, and the Nyquist frequency beyond which the servo theoretically can't see anything, 5 kHz. Aliasing happens, and we see it, but I'll ignore it for now.

Here's where I get out of my depth. There's a set of relationships between the Nyquist frequency, the servo "crossover" frequency, the mechanical stiffness of the actuator (the main frequency is usually around 4 kHz to 5 kHz) and the precision with which the servo system can follow the track centerline. I don't understand all that well enough to explain it very clearly even to myself; I just accept what the servo people tell me. They do the same with respect to what's possible in the mechanics, so I guess we're even!

I guess I'll just have to let it sit here. As the actuator arrives on track, it wobbles back and forth until the servo system is satisfied that it isn't going to be more than 15% (or so) off track on the next wobble, then "seek complete" is posted and it's OK to read data. After the wobbling has dissipated some more and there have been a "sufficient" number of samples within 10% (or so) of the track centerline, the servo system posts "OK to write". You can be sure that a lot of agony goes into deciding what the minimum "sufficient" number is and what the (or so) offset from 10% can be! Anyway, all that wobbling around and decision making is rolled into "settling time".

OK, how do seek time, settle time and latency play together? And how about head switch time, while we're at it?

A hard disk drive is usually sitting on whatever track it last read from or wrote to, and the active head is whatever head was active on the last operation. When the system requires more data to be read or written, the drive hardware translates that into a physical track #, head and sector #. If the data (to read) is already in the drive's on-board cache, it just spits it out, and the hardware is unaffected.

For any write operation, and for any read operation that isn't in cache, these are the possibilities:

1: The destination is a sector on the present head and track.

In this case, the only delay is to wait until the destination sector passes under the head. That's latency. If you assume that the target could be anywhere on the track with equal probability, then on average it'll take half a revolution of the disk to get it under the head (anywhere from just coming up to Shucks! Just went by!)

Latency table:

RPM	Rotation time	Average Latency
3600	16.67 msec/rev	8.33 msec
4500	13.33	6.67
5400	11.11	5.56
7200	8.33	4.16
10000	6.00	3.00

2: The destination is on the present track, but on a different head.

You have to switch heads, then wait the latency period. Actually switching the active head only takes a few microseconds, but it takes some time to make sure that the new head is on the center of the track. Head switches and 1-track seeks usually take about the same amount of time; you have to wait for the next sector to see where you are, and a few more to be sure you're where you want to be. The spec sheets usually say 1 msec or 1.5 msec head switch time, and the average time is usually half that or less. If your spokes come up about every 100 usec, it's about 4 to 6 spokes to

complete a head switch or a 1-track seek.

3: The destination is on a different track, but the same head.

4: The destination is on a different head and track.

You have to seek. Seek time is usually given as an average, and settling is usually included in the seek time. The minimum is a 1-track seek as described above. The maximum is usually in the spec sheet, maybe 18 to 22 msec on a new, fast drive. The time to access the data is the seek/settle time plus latency. A "read seek" is what's usually in the broadsheet. A "write seek" takes a little longer, 0.5 msec to 1 msec longer, because of the additional settling before you can allow writing. Notice, that the additional settling for writing can overlap the latency. On the other hand, if your target comes by during the settle time, then you still have to wait (a whole rev!), so the whole thing cancels out, on average.

On a 7200 RPM, 8 msec seek time drive, the average time to access any sector is about 12 msec. If you knock a msec out of any of the factors--seek time, settle time or latency, you knock a msec off the total. As usual in engineering jobs, you need to balance other factors--acoustics, power consumption, cost of hardware, and so forth, to make a good decision.

=====
Disk Layout - Data Sectors and Servo Sectors
=====

[Note by JT: This is my response to a question in the Compuserve Benchmark forum. I had been conversing about my favorite topic (disk drives) and mentioning servo sectors and data sectors. The writer quite properly asked me to clarify my terms.]

The confusion arises, I think, from two causes. First, the terms in day-to-day use in the industry don't make it explicit what kind of "sector" we're talking about, that is, whether it's a servo sector or a data sector. Second, when one is designing the mechanics of a disk drive, one doesn't worry much about the data; the mechanical engineers design the hardware, the servo engineers design the servo, and the read/write engineers handle the data. We only talk when we have a cross-discipline problem. Of course, that means only once or twice a day! But it turns out that there's more mechanical-to-servo interaction than mechanical-to-read/write.

You know that disks are laid out with concentric cylinders or tracks of data, and that the data isn't arranged in a spiral as it is on a CD. The surface of a current hard disk that uses "embedded servo" or "sector servo" technology has two kinds of data recorded on it--servo information and user data. Both the servo information and the user information is written in discrete units, sometimes called "sectors" and sometimes called "blocks". If one isn't careful to refer to "servo sectors" and "data sectors" explicitly, one creates confusion, as I did.

The servo data is written only once, at the factory, and never re-written afterward. The user data has two sub-categories. I don't know the formal names for them (or even if formal names exist), so I'll call them "format data" and "variable data". The format data is generally written only once, though it is possible to rewrite parts of it. The variable data includes everything that's written when you write a file to disk, so it includes ECC data and FAT (or equivalent) data in addition to your own data.

A disk drive has a fixed number of servo sectors per track. That number is constant across all tracks and across all surfaces of the drive. Furthermore the servo sectors are written synchronously on the disk. The beginning of each servo sector on each track is physically adjacent to the beginning of the corresponding servo sector on each adjacent track. The servo system uses the differences in the recorded patterns on adjacent servo tracks to define the track centerline.

A current "zone-bit recording" or "variable data rate" disk drive has a variable number of data sectors per track. This is different from the old MFM and RLL drives that had all the user data recorded at a single data rate. Variable data rate recording was implemented to get more data on the disk. At the ID of a disk, the track length is roughly half what the length is at the OD, so if you record at a constant data rate, thereby getting a constant number of bits of data per revolution, you record at about twice the linear bit density at the ID as you do at the OD. If you can record twice as fast at the OD, you get the same bit density and twice as much data per revolution.

Data sectors are currently defined with a fixed size, usually containing 512 bytes of true user data. The sector itself usually contains 600+ bytes of data with all the overhead stuff, but it is always a fixed size. So if you record at a variable data rate, you get a variable number of data sectors per track. Current technology does require that a complete data sector exist on a single track, though it is possible to split a data sector across a servo sector.

You can see that one could theoretically have a different data rate for each of the 6200 tracks on a Viking drive, but it turns out to be convenient to divide the surface into 16 zones. Each zone has a constant data rate and a constant number of data sectors. The number of tracks per zone has to vary a bit to make all the numbers come out even, but most zones have the same number of tracks, again for convenience. That's a very incomplete discussion of disk layout, but maybe it will give you some insight.

>> Using the range for the Viking you suggest of 150 to 240, I'm of the opinion that what happens is that the first data zone might have 240 sectors [or spokes or servo wedges] per track [SPT], the next data zone might have 220 SPT, the next might have 190 SPT, ... and the last would have the 150 SPT. <<

The first data zone might have 240 data sectors and does have 72 servo sectors (or spokes or wedges) per track. The next, as you say, might have 220 data sectors, and it has 72 servo sectors. And the last might have the 150 data sectors and the 72 servo sectors.

[Note by AD: In attempting to express the differences between ID and ID-less recording formats I actually eliminated an important feature that John Treder reminded me of how important it was. His comments will follow my ASCII art to show the "big idea" that I removed in the process of editing the ASCII art.]

AD's ASCII art of a single track in a ID-less recording format

```
sector  sector  sector          sector
... [data] [data] [data] [servo] [data] ...
      ^      ^
```

```

      |      |
    gap    gap

```

Previously (traditional) recording format would have IDs where each ID stores track, head, sector and CRC information. The previously mentioned ID-less format increases the amount of data that can be stored.

```

header sector header sector          header sector
... [id] [data] [id] [data] [servo] [id] [data] ...
      ^      ^      ^                  ^
      |      |      |                  |
    gap    gap    gap                  gap

```

[John Treder comments on the "big idea" that Albert Dayes accidentally removed.]

In the course of tinkering with the ASCII art, you simplified out one factor that I think is rather remarkable, and that's the fact that data sectors can actually be split into two pieces, starting on one side of a servo sector and ending on the other. That's true for both ID and ID-less data layouts. And in fact it's the biggest of the non-obvious tricks that makes the combination of sector servo and "constant-density" recording work. When Conner started, they were using constant data rate recording. Every track had the same number of data sectors on it. Therefore they could plan to have an integer number of data sectors between each servo sector. Easy to think about and all that. But you can get (ideally) 50% more data on a disk if you use constant density recording instead of constant data rate. If you use constant density, however, you need to have a variable number of data sectors per track. If you insist on having an integer number of data sectors between servo sectors, you throw away nearly all the advantage of variable data rate. So someone (I don't know who or where) invented a way to "split" data sectors across a servo sector. Suddenly it became practical to put almost the "ideal" number of data sectors on each track, and constant data rate recording became "compatible" with sector-servo.

And that's what I was getting at with my original ASCII art that had:

```

... data  data  da |servo| ta data data data da |servo| ta data ...
      ^^^      ^^^
    split data sector

```

In fact, the original Barracudas had a dedicated servo surface because Seagate's ex-CDC engineers in Minnesota didn't know how to do the split-sector trick. They learned in short order, though!

On Katana (Quantum's internal project name for Quantum Atlas 4 (SCSI) and Fireball +KA (IDE); they're the same except for the interface), there are 124 servo sectors per track, and from 234 to 383 data sectors per track. So there are 1.887 data sectors between each servo sector at the inside and 3.089 at the outside. If you didn't split data sectors, you might throw away 11 data sectors at the outside, it's only about 3%. But at the inside, you'd have to throw away 110 data sectors, or nearly half the track's possible capacity!

If you work things out in complete logical detail, it takes a lot of drawings and thinking, and it works out that you lose a little over 1/3 of the benefit of constant data rate if you don't split data sectors.

=====

GLOSSARY OF TERMS

=====

<g>: grin

<vbg>: very big grin

Actuator: The assembly inside a disk drive that holds the heads and moves them back and forth over the disks.

AGC: Automatic Gain Control - preset for each head and each zone--yet another thing in the tables of data in the "hidden" tracks outside of Track 0. Some disk surfaces are "hotter" than others, and some heads are more sensitive than others. AGC makes it possible to send a fairly consistent signal between the preamp and the main read/write chip on the circuit board. It's convenient to have a stored value for automatically pre-setting the AGC level--then the circuits can fine-tune from there. You could think of it as another scheme to reduce soft error rates.

Aliasing: A type of noise which occurs when a sampled signal contains frequencies which are more than half the sampling device's input sampling frequency. "False frequencies".

ASIC: Application Specific Integrated Circuit

ATA: see IDE

ATAPI: AT Attachment Packet Interface (part of the E-IDE spec).

CD: compact disc - a family of optical storage (e.g. CD (audio), CD-ROM, CD-R, CD-i, CD-R/W, and so forth). These CD members are associated with the "red book", "yellow book", "green book", and "orange book" standards.

CRC: Cyclic Redundancy Check

DSP: Digital Signal Processor

DVD: Digital Versatile Disk

ECC: Error Correcting Code

EE: Electrical Engineer

E-IDE: Known as Fast IDE or ATA-2 and includes support for tape and CD-ROM drives. Created and promoted by Western Digital

EMF: ElectroMotive Force

EMI: ElectroMagnetic Interference

FAT: File Allocation Table

FET: Field Effect Transistor

GB: GigaByte

GMR: Giant MagnetoResistive

ICE: In-Circuit Emulator : A tool (debugging) that emulates the processor on your target system.

ID: header information for a data sector which includes track, sector number, head and CRC information. There is a ID-less recording format which eliminates this header information from the data sector.

IDE: Integrated Drive Electronics defined by Western Digital and Compaq

IEEE: Institute of Electrical and Electronics Engineers

Interface Data Rate: how fast data can be transferred between the host and the disk drive over its interface

MB: MegaByte

MB/s: MegaBytes per second

ME: Mechanical Engineer

Media Data Rate: how fast data can be transferred to and from the media

MFM: Modified Frequency Modulation

MO: Magneto-Optical

MR: MagnetoResistive

MRX: eXtended MagnetoResistive

Nyquist Theorem: The principle explaining the manner in which aliasing is

produced. One half the sampling frequency of any device --
the point above which aliasing is produced -- is often
referred to as the Nyquist frequency of that device.

OEM: Original Equipment Manufacturer.
PCBA: Printed Circuit Board Assembly
PLL: Phase Locked Loop - timing information is in the data stream and is
used to set the receiver's local clock.
PRML: Partial Response, Maximum Likelihood - is the currently popular method
of decoding the analog data that the head reads.
RFI: Radio Frequency Interference
RLL: Run Length Limited
RPM: Revolutions Per Minute
R/W: Read/Write
SCSI: Small Computer System Interface
SI: System International. In SI units the standard length is the meter,
that for time is the second, and that for mass is the kilogram.
SST: Super Sonic Transport (e.g. British Airways' Concorde)
Striping: Storing information on multiple disk drives by splitting up the
information and accessing all of the drives in parallel.
Stroke: The distance across the disk that the heads can move.
TMR: Track MisRegistration
TPI: Tracks Per Inch
UL: Underwriters Laboratories a private company that creates safety
standards and then charges companies to test their products.
VCM: Voice Coil Motor

=====
On-line References : Web Sites
=====

www.1394ta.org - IEEE 1394 Trade Association (Firewire)
www.acm.org - Association of Computing Machinery
www.adaptec.com - manufacturer of scsi host adapters and CD-R/RW software
www.aeanet.org - American Electronics Association
www.ansi.org - American National Standards Institute
www.asme.org - American Society of Mechanical Engineers
www.cardtools.com - hardware/software co-design software manufacturer
www.cmpcmm.com/cc/standards.html - (links to Computer and Communications
Standards)
www.computer.org - IEEE Computer Society
www.disktrend.com - supplier of market research on the worldwide disk drive
industry
www.eg3.com - electronic design, embedded systems, processors, etc.
www.emediapro.net - Emedia Professional magazine
www.fcloop.org - Fibre Channel Community
http://global.ihs.com - Global Engineering Documents
www.i2osig.org - defining a standard interface for high-performance I/O
systems
www.ibm.com - manufacturer of many computer products include storage
devices
www.ieee.org - Institute of Electrical and Electronics Engineers (IEEE)
www.iist.scu.edu - Institute for Information Storage Technology
: Attached to Santa Clara University
www.innovinst.com - Innovative Instrumentation : manufacturer of test equipment
for the magnetic recording industry
www.iso.ch - International Organization for Standardization
www.macdisk.com - Logiciels & Services Duhem
mac/pc data exchange software with links to other scsi

sites

www.mathworks.com - developer of Matlab software

www.ontrack.com - a data (from hard drive and so on) recovery service

www.osta.org - optical storage technology association

www.quantum.com - Quantum : manufacturer of mass storage devices
(of particular interest is the storage resource center)

www.scsita.org - scsi trade association

www.seagate.com - Seagate : manufacturer of mass storage devices

www.storagereview.com - Latest news in the data storage world

www.symbios.com/tl0 - Tl0 Technical Committee (SCSI)
(with links to many other I/O standards, web sites, etc.)

www.tl1.org - Tl1 Technical Committee (IPI, HIPPI, Fibre Channel
standards)

www.tl3.org - Tl3 Technical Committee (IDE/ATA & ATAPI)

www.tully.com/dvdsig - DVD : Special Interest Group

http://yara.ecn.purdue.edu/~smag/ - IEEE Magnetics Society

www.wdc.com - Western Digital : manufacturer of mass storage devices

=====
On-line References : Newsgroups
=====

comp.arch.storage
comp.periphs.scsi

=====
References : (books and magazines)
=====

Spencer W. Ng, "Advances in Disk Technology: Performance Issues",
Computer, May 1998, pages 75-81.

Advances In Information Storage Systems, Volume 8
ISBN: 9810233485

Dictionary of PC Hardware and Data Communication Terms
ISBN: 1-56592-158-5

Digital Magnetic Recording, 2nd edition
ISBN: 1575240777

EMedia Professional magazine
(covers DVD, MO, CD, CD-ROM, CD-R/RW and other optical technologies)
www.emediapro.net

Shen, I. Y. and Low, M., 1996, "Experimental Modal Analysis of a Hard Disk
Drive Spindle Carrying Multiple Disks," Proceedings of the 6th International
Symposium on Transport Phenomena and Dynamics of Rotating Machinery, Vol. 1,
1996

File structures:an object oriented approach with C++, 3rd Edition
ISBN: 0-201-87401-6
(in particular Chapter 3, secondary storage and system software)

Magnetic Disk Drive Technology: Heads, Media, Channel, Interfaces and
Integration

Magnetic Domains : The Analysis of Magnetic Microstructures
ISBN: 3540641084

Microsoft Press Computer Dictionary, 3rd Edition
ISBN: 1-57231-446-X

Shen, I. Y. and Ku, C. -P. Roger, 1995, "On the Vibration Analysis of Multiple rotating Flexible Disks," 1995 ASME International Congress and Exposition

SCSI: Understanding the Small Computer System Interface
ISBN: 0-13-796855-8

The Programmer's Guide to SCSI
ISBN 0-201-18538-5

Upgrading and Repairing PCs, 10th edition
ISBN:0-7897-1636-4
(in particular Chapter 12 on Magnetic Storage)

=====

Additional References : (software or hardware documentation)

=====

WinSCSI-32 by Sequoia Advanced Technology
(Windows 32-bit scsi library - excellent SCSI documentation)
www.seqadvtech.com

$$= \sim = \sim = \sim =$$
[illegible]

Sega Bets It All On Dreamcast!
And much more!

[illegible]

Four of Electronic Arts-Tiburon's Veteran Graphic Artists/Animators, Along With One Former EA Employee, Have Left to Form Their Own Post Production/Graphic Arts Company Called The Ballistic Pixel Lab, Inc.

Four of Electronic Arts/Tiburon's top veteran artists have left EA to form a new company called The Ballistic Pixel Lab.

Donnie Worley, Greg Jobs, Eric Kohler and Jason Hayes have announced their resignation from Electronic Arts/Tiburon and have joined forces with recently departed EA alumni Dave Mathieu to form The Ballistic Pixel Lab. According to company president Dave Mathieu, The Ballistic Pixel Lab will offer cutting edge post production services to the electronic entertainment and advertising industries. The company will offer 3D modeling and animation, 2D graphics, special effects, cinematics development, video editing, logo design, story boarding and web design & development services.

Collectively and individually these skilled artists have played key roles in the graphic and cinematics development of some of EA's most successful and highest grossing game titles, including Madden 97, 98, 99, NCAA Football 98, 99, March Madness 98, as well as other titles in development.

Their work has graced the covers and pages of many of the gaming industries premier trade publications, including Intelligent Gamer, PSX Magazine, Game Informer, EGM (Electronic Gaming Monthly) and Nintendo Power, as well as other marketing devices published by Electronic Arts.

Mathieu says: "The Ballistic Pixel Lab will be the type of Post Production company, by which all future companies will be measured. Our primary goal is to provide our clients the very latest technology has to offer in the field of post production and to deliver this on time with a new standard of quality." While collectively this team has a proven track record in the gaming industry, Ballistic is confident that by remaining proactive they will attain the same level of success in a variety of other markets.

The Ballistic Pixel Lab will soon be moving into a new state of the art facility based in the Orlando Florida area. The Ballistic Pixel lab will be looking for additional funding and any interested parties are welcome to send an email to: invest@ballisticpixel.com or contact Dave Mathieu at (407) 324-0164.

Nolan Bushnell Named to Wave Systems Board

Wave Systems Corp. Wednesday named Nolan Bushnell, 56, electronic gaming pioneer, noted entrepreneur and Chairman and CEO of uWink.com inc., to its Board of Directors. Mr. Bushnell's appointment increases Wave Systems'

Board to seven members.

Mr. Bushnell is best known for founding Atari Corp. in 1971 and serving as its CEO until 1978. Atari, which developed such legendary games as PONG, Breakout and Asteroids, was an industry pioneer and the dominant manufacturer of video games for both consumer and coin-operated markets. In 1981, Mr. Bushnell went on to build Catalyst Technologies which served as an incubator for over 12 technology-oriented companies. Among them were Etak, a creator and manufacturer of automobile navigation systems that was eventually acquired by News Corp., and ByVideo, which developed the first point-of-purchase video kiosk utilizing touch screen technology. From 1991 to 1995 he served as the Chairman of OCTuS, Inc., a public entity that produced totally integrated voice and data communication systems. Mr. Bushnell has also served as a consultant to a variety of companies including Bally Manufacturing, IBM and Commodore International and lectures frequently at colleges and universities throughout the United States as well as at a variety of industry symposia.

Mr. Bushnell currently serves as Chairman and CEO of uWink.com inc., a developer of an e-commerce, games and entertainment network housed on touch screen terminals located in public venues such as airports, bars, train stations, hotels and restaurants. The uWink.com network offers an innovative range of easy to play arcade games, chat and other entertainment features, complemented by a variety of e-commerce capabilities including hotel and airplane reservations. A holder of a variety of patents, Mr. Bushnell is a graduate of the University of Utah, where he earned his Bachelor's in Electrical Engineering. He also attended Stanford University's Graduate School.

''Nolan and I met over 20 years ago as guests on a television show where we were trying to predict the electronic future," commented Peter Sprague, Chairman of Wave Systems Corp. ''We have since worked separately and together to try to make some aspects of that future occur. I have enjoyed his friendship, his vision and his creative enthusiasm. I look forward to working together again. Nolan has always impacted the world around him. Wave will provide him with additional opportunities to reach out and creatively touch the lives of all of us."

Mr. Bushnell commented, ''I'm honored to join Wave Systems' Board as Wave is one of a handful of companies that truly offers technology, security and solutions that can revolutionize commerce and strengthen the world economy. Of particular relevance to my background and focus is Wave's ability to support micro-transactions that enable the reality of the 'pay-for-a-play' arcade to come to the Internet. On a broader scale, Wave supports the sale of wonderful content that previously had no viable economic model, thereby allowing entrepreneurship to flourish on a broad and international scale."

bleem, LLC Announces First Retail Distribution of bleem! Emulation Tool

bleem, LLC, announces that their highly anticipated bleem! software has arrived on retail shelves. With bleem!, PC users can play hundreds of PlayStation games on their PCs, and take advantage of newer 3D hardware to play games at next-generation resolutions. bleem! is distributed by Navarre Corporation, and is now for sale at outlets nationwide, including CompUSA, Fry's, Electronics Boutique, Hastings, and Virgin Megastores.

bleem! has already sold nearly 20,000 copies in just a few months through

www.bleem.com, and has developed a very large following among the gaming community.

'bleem! is truly revolutionary, and we're thrilled to have distributor and retail support," said David Herpolsheimer, President/CEO of bleem, LLC, 'Now tens of millions of PC gamers can play hundreds of great PlayStation games -- most of which have never been available for the PC. And on PCs with recent 3D accelerators, many games actually look better than ever before, with higher resolutions and much more detail."

bleem! sells for \$29.95 in a bright yellow box at retail, and runs in Windows 95 or 98 with DirectX 6. bleem!'s minimum requirements are: a Pentium 166Mhz MMX or higher, 16Mb RAM, high-color (16-bit or higher) graphics card, 2X CD-ROM, Windows-compatible sound card, and 3 Mb hard disk space available.

Discreet and Sony Computer Entertainment Inc. to Accelerate and Enhance Development of Next Generation PlayStation Games

As part of its continued efforts to help game developers reduce development time for new titles, Discreet, a division of Autodesk, Inc., announced its Middleware relationship with Sony Computer Entertainment. As a Middleware partner, Discreet will provide Next Generation PlayStation developers with open access to key character animation technology found in Discreet's recently announced 3D Studio MAX R3 product. Open access to core character animation source code and algorithms in 3D Studio MAX R3 software gives game developers the freedom to integrate leading-edge character animation effects typically seen in broadcast and film directly into their Next Generation PlayStation game engines -- at no additional fee.

'By working closely with Sony Computer Entertainment Inc. to understand the capabilities of the Next Generation PlayStation, Discreet is delivering on its promise to provide game developers with the most-advanced game-authoring environment available," said Jim Guerard, vice president of marketing and product management at Discreet. 'The Next Generation PlayStation is an exciting showcase for 3D Studio MAX R3 software's new surface modeling, skin deformation and soft-body dynamics algorithms. We believe that open access to these character animation technologies is a requirement for leading-edge game developers who want to provide seamless integration between their game authoring and runtime environments."

Discreet recently articulated its strategy for the growing game industry and demonstrated 3D Studio MAX R3 software's character animation technology both as an artist's tool and running in a commercially available game engine. The demonstration used source code components that will be available with 3D Studio MAX and from the Autodesk Developer Network (ADN) -- Discreet's Developer Support Network for game programmers (www.ktx.com/developer/).

'Discreet will promote rapid development of tools specifically designed to let game developers fully leverage the strengths of the Next Generation PlayStation," said Shinichi Okamoto, senior vice president of the Research and Development Division, Sony Computer Entertainment. 'Due to its open architecture and developer support network, 3D Studio MAX represents a powerful catalyst for innovation on the Next Generation PlayStation. Discreet software will create even more ways for game developers to unleash the power of the Next Generation PlayStation."

Discreet's 3D Studio MAX product is a production-proven solution used by professional game developers around the world to create top-selling PlayStation games like Tomb Raider I, II, and III, Madden NFL, Jet Moto III, Ray Man, and Metal Gear Solid. The recently announced 3D Studio MAX R3 software delivers new levels of creative power to game character artists through a host of modeling and animation enhancements that focus on next-generation real-time surface rendering methods. In addition, 3D Studio MAX R3 software's extensive user-interface customization options, MaxScript ``scriptable" plug-ins, and external asset references offer game level builders tremendous flexibility in creating an authoring environment that meshes with their game engine.

Sega Bets Big On Dreamcast

On the same day its parent company's stock was downgraded to near-junk status, Sega of America Inc. announced it had kicked off of its own \$100 million marketing campaign.

"Sega's launch will be the biggest video game launch yet," said Bernie Stolar, president and chief operating officer of Sega of America. The marketing onslaught includes more than 110 TV commercials per week on channels including MTV, and a focus on pre-order sales, which already top 100,000.

"We expect to reach 200,000 by our (September) launch," Stolar said.

Sega's biggest bet is a two-month rental preview of the Dreamcast system. On Tuesday, the game console maker announced a deal with Hollywood Entertainment Corp., the owner of the Hollywood Video rental chain. To prime the buyers' market, 1,000 Hollywood Video stores will begin renting Dreamcast systems on July 15, almost two months before the official launch date.

Eric Lampel of NPD Group thinks the pre-launch rental strategy makes sense. "It will create word of mouth, and if (Dreamcast) lives up to players expectations, it will create demand," he said.

That's good news for Sega, which needs a Western makeover, after being pummeled in Japan for poor sales of its next-generation game console.

According to a Bloomberg News report, Sega of America's parent company, Sega Enterprises, lost 11 percent of its stock value in the Japanese market on Tuesday after Japan's Rating and Investment Information Inc. cut Sega's long-term debt three full notches from "A-" status to "BBB-.". That puts the stock on the edge of "junk" status ranking.

The report blames lackluster sales of Sega's Dreamcast system and its software as the reason for the downgrade. Other analysts have also predicted doom and gloom for the struggling game hardware company.

Yet, while Japanese Dreamcast sales have struggled to hit 1 million six months after Sega's original target date, the United States is a different market, NPD's Lampel said.

Sega of America will eventually make more than 10,000 systems available, said Jeffrey Yapp, president of Hollywood Entertainment.

The roll-out, however, will be gradual. At first, only a single game -- Sonic Adventure -- will ship. Every two weeks, another game will hit the shelves; first will be Capcom's PowerSton, followed by Midway's Ready 2 Rumble and Sega's football title.

At its Sept. 9 launch, Sega of America hopes to have 16 titles available and 200,000 pre-ordered devices. "That's a great size at launch," Lampel said. "A wide variety of titles always helps; they are hitting most of the genres."

Compared to Nintendo Ltd.'s N64 launch in 1996, that's glowing praise. The Nintendo platform wallowed at eight games for its first Christmas season. Sega hopes to have between 30 and 40 this year.

Nintendo turned itself around. Sega is hoping it can do the same.

Hasbro Interactive Celebrates Atari With Release of Atari Arcade Hits: Volume I

CD Compilation Features Original 2-D Classics Plus Atari Historical Archives

Leading entertainment software publisher Hasbro Interactive is proud to introduce the definitive Atari compilation--Atari Arcade Hits: Volume 1 CD-ROM. Atari Arcade Hits includes authentic versions of six of the most popular arcade classics: Asteroids, Centipede, Missile Command, Tempest, Super Breakout and Pong, plus interviews with Atari founder Nolan Bushnell and other memorabilia, all in time for Atari's 27th anniversary.

"Atari games laid the foundation for our booming computer and video game industry today with their fun and addictive game play," said Hasbro Interactive president Tom Dusenberry. "We're proud to bring together some of the most popular Atari games and memorabilia from yesteryear to offer enjoyment to both new and old Atari fans."

In Atari Arcade Hits, first in a three-part series, players relive the fun and excitement of their favorite arcade games of yesteryear, or discover these compelling games for the first time! Play the most popular Atari classics from the '70s and '80s--Asteroids, Centipede, Missile Command, Tempest, Super Breakout and Pong--in all their original glory. In addition to interviews with Atari founder Nolan Bushnell, Atari Arcade Hits includes footage of the original games, as well as flyers, posters and other historical materials. Also included is a collection of desktop themes such as mouse pointers, system icons, fonts, screen savers and more to let Atari fans personalize their computers. Finally, the game contains previews of new and exciting 3D revivals of classic Atari games from Hasbro Interactive. Atari Arcade Hits challenges players of all ages and skill levels and even lets players post high scores on the Internet to compare with arcade lovers from around the globe.

Atari Arcade Hits is available in stores at a suggested retail price of \$29.95. For more information, or to order online, please visit www.hasbro-interactive.com.

'Bugs Bunny Lost in Time' Video Game Adventure Now Available

What's up Doc? Bugs Bunny, the world's lovable hare, now has a video game adventure for Looney Tunes fans of all ages in ''Bugs Bunny Lost in Time," available nationwide now for Sony PlayStation.

Primed to uphold the famed tradition adored by nostalgic fans everywhere and win over a new generation of enthusiasts, the pop culture icon and his legendary cast of cartoon co-stars are at it again, pulling pranks and driving each other ''Looney Tunes" in a test of time.

Bugs Bunny's first 3-D interactive game, ''Bugs Bunny Lost in Time" begins with the harried hare taking a wrong turn in Albuquerque, where he mistakenly activates a time machine, propelling players on a historical crash course.

Retaining the irresistible wit and charm of the classic cartoons, Bugs encounters old nemeses, including Elmer Fudd(TM), Yosemite Sam(TM) and Marvin the Martian(TM) in settings ranging from the Stone Age and Pirate eras all the way forward to Dimension X. With the gamer playing Bugs, the action-adventure title romps through 22 levels of mayhem, boasting traditional ''Tunes" favorites including dropping anvils, bull fights and instant Martians.

''Infogrames has stayed true to the Looney Tunes spirit with authentic-looking characters and backdrops, real voices and a slapstick escapade that will hit home with any nostalgic fan," said Michael Harkavy, vice president, Warner Bros. Interactive Entertainment Worldwide Publishing and Kids' WB! Music.

''Infogrames has offered a new Bugs Bunny experience, combining the antics and adventure of a classic cartoon with one of the most popular video game systems in the world, the Sony PlayStation."

Additional levels pit Bugs against old rivals including Rocky and Mugsy in a 1930s segment, and Witch Hazel, who is up to her old tricks, in a Medieval-themed adventure. The single-player title also offers a tutorial level for beginners with Merlin the Wizard and four bonus levels, which star Bull and Bugs' fiendish friend Daffy Duck.

Incorporating the authentic voices of the current Warner Bros. team and the physical mannerisms of each character, the game is true to the look and feel of the classic episodes, still popular in movie theaters and on television today.

''''Bugs Bunny Lost in Time' will have a broad base of appeal, including fans first and foremost," said Jeff Blanc, director of marketing, Looney Tunes, Infogrames North America. ''Keep in mind that the tales have always been written by adults and for adults, yet the nature of how the final story comes together also hits home with kids right away. The same is true for the game: it will appeal to many audiences."

The title represents Infogrames' first PlayStation product in the company's line of ''Games for the Next Mil-Looney-Um." Bugs and his legendary Looney Tunes cohorts star in ''Carrot Crazy" and ''TWOUBLE," which are both available for Nintendo's Game Boy Color. Infogrames and Warner Bros. Interactive Entertainment have signed an agreement for additional Looney Tunes titles.

'Bugs Bunny Lost in Time" is now available at retail outlets nationwide including, but not limited to, Toys 'R' Us, Electronics Boutique, Babbages, Best Buy, Wal-Mart, Target, Kmart, Sears and Kay-Bee Toy Stores, at an estimated retail price of \$39.95.

Infogrames Marches 'ANTZ' Onto the Game Boy Color

Infogrames, in conjunction with DreamWorks, will release "'ANTZ" for the Game Boy Color in the United States this fall.

Based on the hit animated film, the game will star the famous ant Z, an insecure little bug who has high hopes of finding Insectopia, a mystical world where food is plentiful and all insects live in peace.

''ANTZ' allows kids to jump into the action and act out one of last year's movie blockbuster hits," said Bonnie Scott, product marketing manager, Infogrames. ''We are very excited to be working with DreamWorks. Through this title kids can explore life through a bug's eyes and get involved in a game that has all the charm and humor of the movie."

Ant-oriented hazards presented to the miniature heroes include acidic termites, out-of-control water droplets, gum-soled shoes and, scariest of all, the sun-filled, magnifying glass.

CLCE and Light and Shadow Productions produced and developed the 19-level, movie-based title.

Disney's 'Tarzan' for Game Boy Color System Swings Onto North American Retail Shelves Week of June 28, 1999

Coinciding with the launch of Walt Disney Pictures' animated feature film "'Tarzan", Activision, Inc., in collaboration with Disney Interactive, Inc. , brings Disney's Tarzan to the Nintendo Game Boy Color system. Delivering the action and adventure of the jungle to the palm of gamers' hands, Disney's "'Tarzan" will swing into more than 10,000 retail outlets throughout the United States and Canada the week of June 28, 1999. The game will carry a suggested retail price of \$29.95.

Based on Edgar Rice Burroughs' classic tale, Disney's "'Tarzan" for the Game Boy Color platform challenges players to use their strength, agility and skill to outwit and outlast powerful natural enemies, including crocodiles and packs of wild baboons. Players take on the role of three different characters from the movie, including Tarzan (as a child and as an adult), Jane and Terk, making their way through multiple levels in a mission to save an ape family from an evil hunter.

''Activision is excited to be a part of the launch of Disney's "Tarzan" feature film and Game Boy Color game," said Mitch Lasky, executive vice president, Activision Studios. ''Our Game Boy Color title allows players to bring the action of the Disney movie alive in a pocket sized fun-filled adventure."

Developed by Digital Eclipse, Disney's "'Tarzan" sets new standards for animation and gameplay using fluid character movements adapted from the

Disney film. Artists meticulously re-created each character in the game, allowing players to experience the fluidity and speed on their Game Boy Color. Featuring compelling characters and an intriguing story-line, Disney's "Tarzan" immerses players in an action-packed expedition through a detailed world teeming with challenges. Players must thrash their way through 23 levels, including elephant stampedes and leopard attacks, set in the heart of the African jungle and brought to life with vivid graphics.

Take-Two Interactive Software Inc.'s Rockstar Games To Launch Grand Theft Auto Community on the Internet

Rockstar Games, the high-end console division of Take-Two Interactive Software, Inc., announced that on July 1, it will launch its GTA community world wide web site, www.GTA2.com. Coinciding with the launch of www.GTA2.com. Rockstar will offer exclusively via download from the new site a free expansion pack for GTA:London 1969, GTA:London 1961. Rockstar plans to aggressively promote the site through an Internet advertising campaign which will also commence on July 1.

Rockstar will position www.GTA2.com as the official Internet center for Grand Theft Auto fans around the globe. Along with GTA:London 1961, the GTA the movie short, and release information and links to other GTA fan sites will immediately be available. Over the coming months Rockstar plans to enhance www.GTA2.com with e-commerce features, new demo versions of Grand Theft Auto products, critical release and pre-release information and game hints. At its height the original GTA web site, www.GTA.com hosted more than 500,000 individual user visits per month.

GTA:London 1961 is a complete new level and a valuable prequel to GTA:London 1961. It will be completely free, but will only work if you own a copy of GTA:London 1969, for the PC. GTA:London 1969, which reached the No.1 position in the U.K.'s ChartTrack sales charts, is the first ever mission pack for Rockstar's best-selling Grand Theft Auto. Grand Theft Auto products have shipped more than 2.5 million copies to retailers worldwide in the past eighteen months.

Sam Houser, President of Rockstar Games, said "July 1st is a big day for fans of Grand Theft Auto. Not only are we giving away a free expansion pack, containing some very impressive and fun gameplay innovations, but we are also launching our official GTA2 website, www.GTA2.com. Both events should keep the game's large Internet fanbase very excited prior to the launch of GTA2 this fall."

$$= \sim = \sim = \sim =$$

```
->A-ONE Gaming Online      -      Online Users Growl & Purr!
   _ _ _ _ _
```

Songbird Press Release

SONGBIRD PRODUCTIONS TO ATTEND CGE, ANNOUNCES NEW LYNX GAMES

June 27, 1999

For immediate release:

ROCHESTER, MN -- Songbird Productions recently committed to attending Classic Gaming Expo (CGE), which will be held in Las Vegas, Nevada, on August 14-15. CGE is the premiere event in 1999 for fans of classic gaming systems such as Atari, Intellivision, Vectrex, and more.

"John and Keita have put a lot of effort into making this year's event bigger than World of Atari was last year", commented Carl Forhan, owner of Songbird Productions. "I'm thrilled that Songbird will be a part of CGE, and I plan on having a variety of merchandise for sale for many Atari systems."

Additionally, Songbird Productions hopes to keep the momentum growing for the orphaned Atari systems Lynx and Jaguar. To that end, two new Lynx products will be launched at CGE.

The first is Lexis, an interesting twist on the "falling blocks" concept familiar to gamers everywhere. In this game, you don't have to clear a row of bricks -- you have to form words out of falling letters. This fun and challenging game includes a dictionary of 20,000+ words, some of which are special cheat words recognized by the game. Several play modes are also included.

The second Lynx product is "Crystal Mines II: Lost Caverns". Songbird Productions and Serious Cybernetics have teamed up with original Crystal Mines I and II developer Ken Beckett to take advantage of a hidden feature in the existing Lynx Crystal Mines II cartridge -- it supports downloads of new levels into Lynx RAM via the comlynx port. The product will include Win95 software on a CD-R and a special Lynx-to-PC serial cable which also doubles as a developer cable for those who may own a BLL or SIMIS cartridge for the Lynx.

CGE promoter John Hardie recently expressed his enthusiasm with regards to the news, saying, "We're delighted that Songbird Productions, the premiere Lynx and Jaguar developer, has chosen CGE '99 as the venue in which to debut and display their new products."

Songbird Productions also anticipates being able to demo unreleased Jaguar games at CGE, including Protector and Skyhammer. At press time, it was uncertain if these games will be available for purchase in time for CGE.

Songbird owner Carl Forhan recognizes that he's not alone in his endeavor to support Atari platforms. "Several sponsors have helped make it possible for Songbird to attend CGE '99. Ken Beckett and Multimedia 1.0 each deserve a big thank you for their generous sponsorship of Songbird Productions."

To keep up to date with the latest news at Songbird Productions, be sure to visit the company web site at <http://songbird.atari.org>, or send an email to songbird@atari.org.

Copyright 1999 Songbird Productions. All rights reserved. This article may be reprinted in its entirety.

=~==~==

A-ONE's Headline News
The Latest in Computer Technology News
Compiled by: Dana P. Jacobson

Experts See No Microsoft Antitrust Settlement Soon

After months of trial that exposed Microsoft Corp.'s bullying business practices and cast doubt on the credibility of chairman Bill Gates, the U.S. government is unlikely to settle the case unless the software giant makes major concessions, experts say.

The government has little incentive to compromise because most experts believe District Judge Thomas Penfield Jackson will rule against the software giant.

"There's no question the government has made a case and I would be very surprised if Judge Jackson didn't come down with a finding to that effect," said Thomas Morgan, who teaches antitrust law at Brigham Young University in Provo, Utah.

Another antitrust lawyer said it is too soon for either side to settle.

"Right now, hope springs eternal for both sides," said Kevin Arquit, an antitrust lawyer for Rogers and Wells in New York. "Microsoft and the government are starting from different ends."

The Justice Department and 19 states allege that Microsoft leveraged monopoly power derived from its Windows operating system by bundling its Web browser to hurt a competing product made by Netscape Communications.

But Gates said the integration of Internet Explorer into Windows was a "fantastic thing" that made the product better.

"I'm surprised we have to defend it at all," he said in Washington this month.

Although he never appeared in person at the trial, Gates was a key figure. On the trial's first day, last Oct. 19, a government lawyer showed direct contradictions between the chairman's videotaped testimony and memos he had written to other executives.

Judge Jackson shook his head in amazement and smiled -- while the courtroom audience laughed out loud -- at some of the more tortured exchanges between an evasive Gates and government lawyer David Boies in tape excerpts played over several weeks.

Boies held out an olive branch to Microsoft on the courthouse steps soon after trial testimony ended last Thursday, giving a nod to Microsoft's lawyers.

''We fight in court sometimes but more than that we cooperate in court," said Boies.

If a settlement is possible, he said, ''nothing that happened in court will get in the way of that."

But what happened out of court a day later did get in the way of settlement.

The Justice Department bristled at a news report Friday that suggested the two sides were negotiating. A spokeswoman blamed Microsoft for violating a March pact to keep such talks a secret.

''Microsoft's repeated disclosures to the press to spin its position in settlement are a significant obstacle to making progress," said Justice Department spokeswoman Gina Talamona.

Microsoft responded in kind with a statement that appeared to blame the Justice Department for the leak, which included an outline of Microsoft's proposals.

The tiff left unclear what would be needed to settle the case. Some experts believe any settlement that the government would agree to would have to stop Microsoft from using monopoly power to crush future competitive threats.

Arquit said a preliminary judicial decision expected in September or October, known as ''findings of fact," may help spur new talks. The findings will point the way to Jackson's later decisions in the case.

Jackson will sift through the evidence -- including the testimony of 26 witnesses and several thousand exhibits such as e-mails, expert reports and videotaped demonstrations -- and then say in writing what he believes is the truth.

''When the judge issues his findings of fact, like it or not, that will create a benchmark from which discussions can flow," said Arquit.

Some critics of Microsoft wonder if Jackson may follow the example of District Judge Harold Green, who used a routine decision to send an unmistakable message to AT&T in its landmark antitrust case.

''The Bell system has violated the antitrust laws in a number of ways over a lengthy period of time," wrote Greene in his Sept. 11, 1981, ruling, warning that from then on ''the burden is on the defendants."

His strong words led directly to AT&T's voluntary decision to break up on Jan. 1, 1984.

Microsoft Says SEC Probing Accounting Practices

Federal authorities are investigating Microsoft Corp.'s practice of setting aside some of its software revenues and recognizing them later, chief financial officer Greg Maffei said Wednesday.

Maffei said in a conference call with reporters and analysts that the U.S. Securities and Exchange Commission apparently launched the probe after a

newspaper report drew attention to a dispute over whether the company dips into its "cookie jar" of reserves to smooth financial results.

"We don't know the entire scope other than it relates to reserves and reserve policies," Maffei said, adding that the company is "cooperating fully" with the SEC.

The brief Wall Street Journal article, published in January, referred to a lawsuit brought a former company auditor who charged he was wrongfully terminated after raising questions about so-called "cookie jar" accounting.

Microsoft, well known for its conservative financial reporting practices, has consistently said it follows industry guidelines in deferring recognition of revenue from products such as Windows and Office to account for later delivery of upgrades and support.

Microsoft began the practice with the launch of Windows 95, saying such software was becoming more like magazine subscriptions that are paid for up front and delivered over months or years. Since then such "unearned revenue" on the company's balance sheet has grown to more than \$4 billion.

The SEC declined even to confirm the investigation, although Chairman Arthur Levitt has waged a campaign against improper use of reserves to pad corporate results in later quarters.

Analysts said they were not concerned by the SEC probe.

"I think it's much ado about nothing because you can still see the revenue and track it," said analyst Chris Galvin of Hambrecht and Quist. "It's not like they're holding back."

Bill Epifanio, an analyst with J.P. Morgan, said Microsoft's accounting practices had been consistent and conservative.

"Given the size of their revenues it's turned into quite a valuable asset for them in terms of helping them smooth out their results," he said.

He and others said Microsoft stock could go higher Thursday on separate remarks by Maffei indicating that business in the quarter just ending was at least as strong as expected.

The remarks came as Maffei disclosed that Microsoft was adjusting its reserve policies to comply with new accounting guidelines and recognize more Windows and Office revenue up front.

Maffei said the change was unrelated to the SEC probe and had been discussed with the agency.

The change will add about one cent a share to Microsoft earnings in the fiscal fourth quarter that ended Wednesday and another one cent a share in fiscal 2000, he said.

Maffei also said the company had reclassified revenue and expenses from four relatively small businesses including product support and Internet access.

The change had the effect of adding about \$2 billion in revenues over the past 11 quarters, boosting the total by about 5 percent, although because of related expenses the company's reported net earnings were not affected.

SEC Seen Focusing On Non-Public Microsoft Reserves

With the disclosure that Microsoft Corp. is being investigated by the U.S. Securities and Exchange Commission, the software giant once again finds itself the target of an increasingly aggressive federal agency.

Microsoft chief financial officer Greg Maffei disclosed the SEC investigation late Wednesday, saying only that the agency is examining its ``reserves and reserve policies."

The investigation is far less serious than some of Microsoft's other legal problems including the landmark antitrust suit brought by the Department of Justice and 20 states, but it still could damage the software giant's image even if no action is taken, analysts said Thursday.

``I think any time you have an SEC investigation from a PR perspective it's hard to put a positive spin on it," said Andrew Roskill of Warburg Dillon Read.

Nevertheless analysts said they did not expect the probe to have a major impact, and Microsoft stock rose \$1 to close at \$91.19 in heavy Nasdaq trading.

Maffei, who said Microsoft is cooperating fully with the SEC, said he believed the investigation was triggered by a newspaper article referring to accusations the company used "cookie jar" accounting to smooth out financial results.

The accusations that Microsoft manipulated its cash reserves to provide more orderly earnings were brought by former company auditor Charlie Pancerzewski in a wrongful-termination lawsuit that was settled last year.

SEC Chairman Arthur Levitt has waged a campaign against such techniques, saying companies should resist Wall Street pressure to adjust earnings to meet targets.

``When accounting practices are defined more by their gimmickry instead of their representation of underlying business conditions, the public trust is jeopardized," Levitt said last year.

Dole: Block Pornography On Library PCs

Republican presidential hopeful Elizabeth Dole's call for public libraries to install software filters on all PCs - even those used only by adults - is drawing criticism from some high-tech quarters and applause from anti-pornography crusaders.

Dole, in a stump speech in Bellevue, Wash. Monday, called on Congress to withhold federal Internet access subsidies from public libraries unless they install content filtering software on all computers used by patrons.

"To protect our families and to protect the taxpayers, we shouldn't let pornography gain access to federally funded libraries through an electronic

back door," the former Cabinet secretary and American Red Cross president said in prepared remarks later posted to her Web site.

"Pornography is off-limits to children but it is readily accessible to adults. That is wrong," Dole said. Dole also said that she has sent letters to House and Senate leaders detailing her library filtering proposal.

Some free-speech watchdogs pointed out that, beyond the First Amendment problem of restricting adult access to sexually explicit materials protected under the Constitution, such a move might also block adult library patrons' access to other online material.

"The technologies that have been mentioned in all the bills on this issue are not perfect now and they may never be," said Ari Schwartz, a policy analyst with the Center for Democracy and Technology. "This would put health information at risk, since many filters block health sites. It would be an especially big problem for the poor and people who don't have computers in their homes," Schwartz said.

Well-intentioned but clumsy filtering programs have become legendary for filtering sites that mention potentially naughty words such as breasts-as in "grilled chicken breast" or "breast exam."

Technical glitches notwithstanding, the House this month approved a measure to restrict federal Internet access funds to libraries and schools using the filters on PCs accessed by youngsters. A similar bill is pending in the Senate.

Other observers said Dole's comments were aimed at religious conservatives and added that the high-tech industry is used to playing the heavy on politically sensitive issues such as pornography.

"Overall, I think this industry is bemused that Washington either panders to [people's fears] in the hopes of getting campaign contributions or portrays [pornography] as evil incarnate in order to squeeze out a few cheap votes," said Dave McClure, executive director of the Association of Online Professionals trade group.

Dole's proposal drew cheers from an activist group that was a key supporter of the original Communications Decency Act, which would have made "online indecency" a crime.

The Supreme Court declared the CDA unconstitutionally broad in 1997 and proclaimed the Net worthy of the highest available level of free-speech protection.

Anti-porn activists have nonetheless continued their fight against the threat of online smut. "From the perspective of those of us who have worked on the porn issue for 10 years, I do agree with her," said Monique Nelson, COO of Enough is Enough. "There have been too many incidents of kids getting access to porn sites in public libraries for this issue to continue to be ignored" by politicians, she said.

Since most libraries would not consider stocking videotapes of pornographic movies for fear they would fall into children's hands, they should also deploy Net content filters, if possible, even on machines reserved for adults, Nelson said.

"Even if you have a children's section and an adults section, kids have a way of getting around" such boundaries, she said.

Critics of Nelson's stance often point out that such arguments are subject to easy abuse. In Oklahoma, an ideologically conservative coalition tried to ban the Oscar-winning movie, the "Tin Drum" as pornographic.

So far, at least one court has ruled that library filters violate the First Amendment. Last year a federal court in Loudoun County, Virginia ruled that the software filters blocked Constitutionally protected materials and could not be used in public libraries.

Harvard Caught In Hacker Crossfire

Harvard University is caught in the middle of an online war between hacking-scene follower AntiOnline.com and the hacking community at large.

On Wednesday, the Cambridge, Mass., university removed an independent security Web site, known as Packet Storm, which it had been mirroring on its servers for only 10 days.

The reason: A directory of material hidden in the Web site, and thus on Harvard's servers, that had "sexually related material and personal attacks on an individual not affiliated with the University," said Joe Wrinn, director of news and public affairs for Harvard, in a statement released by Harvard on Thursday.

"We agreed to have a site that had security-related materials only," said Wrinn. "Both parties involved were using us in a way that was completely inappropriate."

Ken Williams, a North Carolina State University employee and the Webmaster of Packet Storm, angrily refuted the allegations.

"This statement is incorrect, and even libelous itself by implying that I had 'sexually related material' on the server," he wrote in an e-mail. "I never did!"

According to Williams, the directory -- labeled "/jp" because it was a collection of material satirizing AntiOnline founder and chief John P. Vranesevich -- had a parody of the AntiOnline site.

But others familiar with the site said that the parody also contained photos of nude women that were intended to be more sarcastic than sexual. Harvard obviously didn't get the joke. Harvard's Wrinn did not know specifically what sort of "sexual" content was contained on the site.

Harvard in the hot seat "We are in the middle of this and it's inappropriate," said Harvard's Wrinn, sounding distinctly uncomfortable with the attention that the issue was attracting. Harvard intends to send the complete contents of the site back to Williams so that he can post it elsewhere.

No wonder: Packet Storm wasn't just a small-time site -- it had been the place to go for both hackers and security experts to get up-to-date security information.

"Packet Storm was a huge compilation of security tools," said Brian Martin, known as "Jericho," the Webmaster of hacker news and information site

Attrition.org. "It was updated daily with tools. It was always there."

Among organizations that used and mirrored the site: The Department of Defense and the Federal Bureau of Investigation, claimed Webmaster Williams.

'I didn't have an anti-J.P. Temple of Hate' Yet, Williams had also sided with many others in hacker circles who have been waging a war -- of mainly -- words against AntiOnline's Vranesevich and his latest ally, Caroline Meinel, security researcher and webmaster of The Happy Hacker.

"I didn't have an anti-J.P. Temple of Hate or anything," said Williams. "But there are companies, organizations, and individuals out there that [we believe] are black-eyes of the industry."

So, Williams attached a non-public directory to the Web site that archived parodies and criticisms of AntiOnline's founder.

The directory represented a single facet of a complex war of image in the hacker not-so-underground. For the most part, AntiOnline and its main foe, Attrition.org, have squared off with conflicting allegations of slander, libel and plagiarism.

Hitting close to home For AntiOnline's Vranesevich, the directory buried inside of the Packet Storm site hit a little too close to home.

"I can understand a parody -- I have no problem with that," said the 20-year-old Pennsylvania Webmaster, adding that he thought Williams crossed the line into poor taste by adding high school yearbook pictures of Vranesevich and his family to the online archive.

Williams acknowledged that the photos had been put up, but that since they had come from a source already online, the Packet Storm Webmaster thought the pictures were fair game.

Vranesevich's answer? The Webmaster notified Harvard of the hidden directory in a letter to the university's provost -- and Harvard quickly took the site down.

Did Harvard act too quickly? B.K. DeLong, a Boston-based computer security consultant, thought Harvard acted too quickly.

"I am kind of disappointed that an institution like Harvard was so quick to pull the plug just to avoid a potential suit," he said.

Yet, Harvard wasn't the only one to act quickly. By late Wednesday night, the Keebler Elves -- the cybergang that claimed responsibility for hacking into the National Oceanic and Atmospheric Administration last week -- defaced another government Web site with the news.

"Now, because [of] JP ... Packetstorm is no more, and never will be again," the hacked site lamented.

Unnamed hackers also struck at AntiOnline more directly. AntiOnline's site came under a denial-of-service attack -- which floods a particular site with random data -- so severe that its Internet service provider pulled the site for almost 12 hours on Thursday, said Vranesevich.

Ugly threats Other attacks were even less friendly. "I have received more death threats in the last 24 hours by phone, than I have in five years," he

said.

Not quite an apology, Vranesevich added that he never intended the entire Packet Storm site to be taken down.

"I know what it's like to have the university stomp its foot down on you. When I was a student at the University of Pittsburgh, I had my Web site shut down," he said. "But I never threatened anyone."

In his mind, the contents of "/"jp" did.

Game Blocker Will Be In Next Windows Version

Parents will be able to restrict the kinds of games their children play by enabling a feature to be included in a new version of Microsoft Corp.'s Windows operating system due out next year.

The Redmond, Wash.-based software company is currently working with game manufacturers to embed code information on ratings in their new games, which the Windows Game Manager would check, said Kevin Bachus, product manager for Microsoft's multimedia group. Game ratings are currently printed on game packaging.

The feature would allow parents to keep games containing excessive violence, sexual content or expletives from being played on the computer. Microsoft's Internet Explorer has a similar feature that can be used to block access to certain Web sites, he added.

Microsoft's planned move comes at a time when the computer and video game industry is under siege because of concerns about the effect of such violence on children and its possible connection to the school shootings around the country.

Douglas Lowenstein, president of Interactive Digital Software Association, the Washington, D.C.-based game publishers' lobbying group, said he thought the feature sounded like a useful tool and predicted support for it from the software industry.

"I think philosophically it's the kind of thing as an industry we've always been supportive of - parental empowerment tools," Lowenstein said. "Anything that facilitates the control of what kids have access to is positive."

Ann Stephens, president of the Reston, Va.-based research firm PC Data Inc., questioned whether the program would be effective.

"I think any technologically savvy kid is much more technologically savvy than their parents. Maybe this will give some parents some peace of mind, but I don't think it will have any impact," Stephens said.

Stephens said a recent PC Data survey found that most people believe violent movies and television shows are more damaging to children than computer games.

"The question is this - do we ban everything because there are crazy people out there and we do not know what sets them off?" Stephens asked. "I think that's overkill."

Bachus said Microsoft was confident the blocking feature could be made easy enough for a non-computer-savvy parent to use as well as secure via a password.

''There will always be ways to skirt around this but ... we think it gives parents another tool to try to be a better parent," Bachus said.

Parents also would have the option to specifically approve certain games or programs, such as a word-processing program, that might not be rated, Bachus said.

AOL To Unveil Version 5.0 This Year

America Online Inc., the world's largest Internet access provider, will soon unveil AOL 5.0, software that promises faster registration, easier setup and enables users to send photos as easily as e-mail, the company said [Tuesday] in a statement.

AOL, based in Dulles, Va., said version 5.0 will begin beta testing within a few weeks and will be distributed to subscribers later this year.

AOL, with more than 19 million subscribers between itself and its CompuServe subsidiary, said the new software includes a new Welcome Screen, with direct access to the AOL Channel Guide, e-mail and top news headlines.

New features include ''You've Got Pictures," developed in partnership with Eastman Kodak, for photo sharing; ''My Calendar," an interactive calendar; and ''My Places," which will allow members to personalize their welcome screens with their favorite sites.

AOL 5.0 will also support e-mail use on Palm electronic organizers. AOL and 3Com Corp. recently reached agreement on a deal to bundle AOL software with Palm organizers.

AOL, Microsoft Offer Free PCs

Major Internet services and computer makers are joining forces to offer a seemingly unbeatable deal to entice millions more people to join the Web revolution: free personal computers.

While the freebies have conditions, a flood of offers this month will hasten the day when PCs are routinely given away just like cellular phones.

America Online, Microsoft and Prodigy, three of the biggest names in Internet access, are promising free personal computers to anyone who agrees to buy their service for three years. Other major Internet providers are expected to add similar deals with computer makers later this month.

The giveaways don't provide a \$100-\$300 monitor and demand a 3-year commitment to the Internet services, which adds up to as much as \$790.

But home computers, which used to cost as much as a used car, are clearly hitting bargain basement levels. While smaller companies started giving

away PCs last year, announcements on Wednesday by AOL and on Thursday by Prodigy show how plunging prices are fast making PCs into just a delivery mechanism for services and software. A steady drop in prices of parts, such as microprocessors and hard-disk drives, has helped make giveaways possible.

''I think it will be nearly universal," said Roger Kay, an industry analyst with the International Data Corp. research firm. ''It's like the old AT&T model, where they gave you the phone."

AOL said it would give a \$400 rebate on selected computers made by eMachines Inc. to anyone who subscribes to AOL's \$21.95-a-month CompuServe Internet access for three years. The lowest-price eMachines computer is \$399, in effect making it free, but the monitor isn't included. The deal also applies to all computers bought at the Circuit City retail chain.

The arrangement is expected to help AOL finally bolster the subscriber base of its CompuServe service, which is targeted to small businesses and home office users, but has only 2 million subscribers. AOL's flagship online service has 17 million subscribers.

Prodigy announced a similar 3-year deal for any brand of computer bought at Best Buy consumer electronics stores. Its Internet access costs \$19.95 a month.

Microsoft Corp. is offering a \$400 rebate to customers who buy a personal computer at Staples Office supply stores. The one-week trial ends Saturday, but Microsoft officials said they are considering extending the offer, as they did with a 3-month old offer at the small computer chain Micro Center.

The major PC makers, including Hewlett-Packard Corp. and Compaq Computer Corp. are expected to announce similar deals with Internet service providers later this month.

The moves will squeeze already thin profit margins at PC makers, which may be forced to offer lower-priced machines to compete with eMachines' budget product.

But they may have little choice. EMachines, a startup in Fremont, Calif. that launched its first major sub-\$600 PCs last September, has vaulted into No. 5 among PCs sold at retail with 9 percent of the market in May, according to the research firm PC Data. Compaq was No. 1, followed by Hewlett-Packard, Apple Computer Corp. and IBM.

A chief reason many consumers don't want to spend a lot is because many use PCs mainly for Internet access.

''People are beginning to look at the Internet expense as a monthly annuity, just as you pay the telephone company," said Stephen Dukker, the chief executive of eMachines.

Thanks to the Internet, people are quickly getting conditioned to free stuff. Nowadays, it's possible to get free Internet access, e-mail, faxes and voice-mail messages sent to your computer for playback over its speakers. Many of these freebies require users to watch advertising or give up personal information to marketers.

Not that the PC offers are without risks. Signing up for a three-year contract is a lifetime in Internet time. New methods of accessing the Internet far faster than traditional telephone modems, such as through

cable TV lines, are becoming available to more Americans. And customers who break their three-year contracts with Internet providers face penalties.

Web Group OKs XML Design Standards

An Internet standards body has approved coding specifications that will allow the designing of Web pages using Extensible Markup Language, or XML.

The World Wide Web Consortium on Tuesday approved a standard for using "style sheets" in conjunction with XML, giving Web designers a wider range of design and display options when using XML.

Style sheets allow documents in an XML format to be easily converted into HTML, or Hypertext Markup Language, which is the programming language used for displaying documents on the Internet.

XML documents can be coded to identify any portion of a document, such a quote from a newsmaker. With this coding, search engines can give users more refined search, finding only those documents where a specified newsmaker is quoted. XML coding also could specify the zip code or codes that a document's contents refer to, allowing searches to find information about a specific zip code.

Microsoft's Internet Explorer 5.0 browser already supports the new coding specifications, while Netscape says future versions of its browser will.

The World Wide Web Consortium was founded in 1994 to develop common standards for Web design. The consortium is jointly hosted by the Massachusetts Institute of Technology, the Institut National de Recherche en Informatique et en Automatique based in France and Keio University in Japan.

AOL Offers Digital Photo Service

All those shoeboxes filled with snapshots may be history.

America Online on Tuesday began offering a service that allows subscribers in Orlando, Tampa and Cleveland to take conventional film to 38,000 Kodak photofinishing stores around the country and have digitized copies delivered to their computer for an extra \$5 to \$7 a roll.

The service, called "You've Got Pictures!" allows subscribers to store photos in a private, online photo album. They are able to personalize the albums by adding captions and backgrounds and cropping or enlarging the pictures. The photos could either be downloaded or transmitted online to friends or family.

The service will be available to all of AOL's 17 million customers by the end of the year, said Debbie Fletter, AOL's director of programming and promotions.

Customers will still be able to pick up their negatives and prints from their photofinishing stores.

Orlando, Tampa and Cleveland were chosen as test cities because they have large numbers of AOL subscribers, Ms. Fletter said.

AltaVista - Not Just A Search Engine

AltaVista Co. today rolled out several new services intended to begin the company's transition from a single site to a network, including search capabilities, portal services and shopping.

Some of the services, like a "freshness guarantee" for AltaVista's search results, are designed to promote the company's reputation for cutting-edge search technology, while others, such as a redesign of Shopping.com and the introduction of "local portal services," incorporate recent corporate acquisitions.

AltaVista president and CEO Rod Schrock said the efforts renew the company's "laser " focus on the growing population of experienced Net users. "We're going to carve off that audience. We're going to own them," he said. "It's going to be harder from now on to have a one-size-fits-all site."

AltaVista, which started off as a showcase for the servers of Digital Corp. (later purchased by Compaq Computer Corp.), has maintained high levels of traffic largely because of the positive reputation of its Internet-searching technology. But the site has struggled to find a business direction to rival such so-called portals as Yahoo! Inc. or Excite Inc.

Compaq is currently in talks with CMGI, a holding company that owns a large stake in Lycos Inc. and may purchase AltaVista as a flagship for its own network of Internet properties.

Only the freshest Modifications to AltaVista.com include a "freshness guarantee," which promises that all pages are less than 28 days old, as well as a multimedia search feature. The search-results pages now include product offers related to the search terms.

The company introduced the MicroPortal, a piece of downloadable software that includes current personalized information and such features as search and, later, instant messaging. The software is an update to a product called AltaVista Tracker, introduced in March for users of Microsoft's Internet Explorer 5.0. To promote MicroPortal, AltaVista is giving away \$5 million in credit at Shopping.com.

Shopping.com itself got a facelift, including a redesign of its seven category-oriented stores and a launch of new customer-service features. The store will also hold promotional "Ultimate Brand Auctions," in which brand-name items are auctioned off with a \$1 minimum bid.

Finally, the company is rolling out "local portal" sites, which combine the features of an Internet hub with local information provided by Zip2 Corp. The first to roll out is Houston, in cooperation with the Houston Chronicle.

Schrock said the company plans to unveil a complete "network" by the end of the year, including 25 new features for AltaVista.com.

Compaq Faces Uncertain Future

The name Compaq Computer Corp. once conjured up images of a well-managed company with a Midas touch, but its golden reputation has been so tarnished that some are beginning to wonder about its future.

For the moment, the questions are less about whether Compaq will survive and more about how the computer giant can reshape itself as a nimble competitor in a fast-changing industry.

While Compaq is still the world's No. 1 personal computer maker, there are those who doubt the Houston-based company will make it in its current form unless it gathers itself in a relatively short time.

"I think they've got about 24 months to put their bets down -- and if they don't get them down on the right numbers in that period, I think the company is probably going to be a candidate for a sellout," said analyst Ernest Widmann of Widmann, Siff & Co.

Since the first quarter of 1998, Compaq has been struggling to find its way through a maze of excess inventory and the controversial \$8.4 billion acquisition of Digital Equipment Corp. The problems cut profits, sent the stock plummeting, and in April led to the dismissal of Chief Executive Officer Eckhard Pfeiffer.

While many welcomed a change at the top, analysts said it did not solve Compaq's key problem -- a production and distribution system that has been left in the dust by competitors such as Dell Computer Corp. and Gateway Inc.

Since Pfeiffer's departure, Compaq Chairman Ben Rosen has led a management team that is simultaneously looking for a new chief executive and trying to reshape the company.

So far, no new leader has been named and, analysts say, the company's strategy is unclear.

Compaq has yet to let the investment world know how it will address the dilemma of competing against the direct sales model used by Dell and Gateway while at the same time holding on to the resellers who have historically been Compaq's primary outlet.

The company has said the Internet is a key part of its future, but Tuesday it sold a majority stake in its AltaVista search engine, which came with Digital, to CMGI Inc. for \$2.3 billion in stock and debt.

"Please don't believe that we're withdrawing from the Internet," Rosen said in a conference call with analysts and reporters. "This is part of a strategy to enhance our position."

The move got lukewarm applause from analysts: They wondered if Compaq was selling off a potential moneymaker but also said the sale would enable management to focus on getting its computer business in order.

"We believe it makes sense for Compaq to hand off its AltaVista division to a company such as CMGI that is better suited to handle an Internet

business," said Steve Milunovich of Merrill Lynch.

But, he added, "We continue to be neutral on shares of Compaq, given the low visibility regarding management and strategy."

Compaq stock gained \$1.06 to \$23.375 on the New York Stock Exchange Tuesday after the AltaVista deal was announced but remains well below its 52-week high of \$51.25. Wednesday, it edged up 31 cents to \$23.69 in composite NYSE trading.

Analyst William Conroy of Sanders Morris Maundy said Compaq's strategic decisions were complicated by changes taking place in the computer industry, which he believes is becoming dominated less by hardware and more by value-added services.

"It's got its own world (of issues) to deal with, but the world in which it plays is also shifting. The sands are moving beneath it," he said.

Part of that change is the simple fact that computer prices are falling and margins are harder to maintain, said Widmann.

"If you're not right on the money with your marketing and your technological position, you're going to get squeezed. And these guys (at Compaq) are one of those being squeezed out right now," he said.

Apple Sues Over iMac Look-alike

Apple Computer Corp. has filed suit against the makers of an iMac look-alike, asking that they be prohibited from distributing the system.

The suit, filed in federal court against Future Power and Daewoo Telecom, complains that the E-Power systems, which are Intel-based, "illegally copy Apple's designs."

"There is a universe of original designs that Future Power and Daewoo could have created for their computers, but instead they chose to copy Apple's designs," Steve Jobs, Apple's interim CEO said in a release. "We've invested a lot of money and effort to create and market our award-winning designs, and we intend to vigorously protect them under the law."

The E-Power systems debuted last month at PC Expo. The all-in-one systems will come in an array of colors -- amethyst, ruby, topaz, emerald, and sapphire -- just like Apple's iMacs. The systems were scheduled to ship in September, starting at \$799 for a system using a 466MHz Celeron processor.

Apple famously lost a suit against Microsoft Corp. in the late 1980s, when it charged the software maker with stealing the "look and feel" of the Mac OS.

In the current suit, the company is asking the court to prohibit Future Power and Daewoo from distributing the systems and award Apple actual and punitive damages.

Future Power is a joint venture with Daewoo Telecom, based in Santa Clara, Calif. A spokesman for Future Power was not available for comment.

Microsoft Acknowledges Windows 98 File Checker Bug

Be careful when checking your system files.

Microsoft Corp. has acknowledged a Windows 98 bug that can render PCs with the operating system unbootable after running the System File Checker tool.

The bug, acknowledged by Microsoft last week, occurs when the checker installs the wrong file in place of a damaged or missing Windows 98 system file. The wrong file can be installed because incorrect information was provided to the checker about the location of "cabinet" files, which is where Windows 98 system files are stored. Because of that incorrect information, the System File Checker may locate and install the wrong file in place of the system file. As a result, PCs can freeze and be unable to load Windows 98.

According to Microsoft, users who experience the problem can work around it by starting their computer with the Windows 98 Startup Disk and replacing the replacement system files using extract.exe at a command prompt.

The bug affects only Windows 98, according to Microsoft.

Microsoft Makes Windows 98 Updates Downloadable

Customers who don't want to pay \$5 for the latest patches and bug fixes to Windows 98 can now download them from Microsoft Corp.'s web site.

Microsoft has made available the individual components which comprise Windows 98 Service Pack 1. These include Windows 98 System Update, Internet Explorer 4.01 Service Pack 2, Windows 98 Year 2000 Update, Windows 98 Year 2000 Update 2 and Outlook Express Year 2000 Update.

Microsoft had committed earlier this year to making the various SP1 fixes available for download from its web site by early summer.

While \$5 may not seem like much to obtain a single CD with all these fixes, some testers chafe at the idea of paying anything for bug fixes, notes Nate Mook, web master for the tester site BetaNews.Com.

"Many of our visitors, along with other people I've talked to, held out not wanting to pay for something that should be offered as free support, especially since most of the updates relate to Y2K issues," says Mook.

Microsoft last month released the CD version of Windows 98 SP1. Earlier this month, the company made available through retail its Windows 98 Second Edition (SE) upgrade product. SE includes all the bug fixes that are part of SP1, plus a handful of new features including Internet Explorer 5; NetMeeting 3 conferencing software; Internet Connection Sharing, home networking technologies that allow multiple home PCs to share files, printers and a single Internet connection; and improved support for Universal Serial Bus, IEEE 1394 and ACPI.

Windows 98 SE sells for an estimated retail price of \$109. The upgrade version--called Windows 98 Second Edition Updates--is available to Windows 98 users only. It is priced at \$19.95 (plus shipping and handling) and

includes a CD with all the fixes and new features, plus a book detailing the use of SE in an Internet environment.

=~==~==

Atari Online News, Etc. is a weekly publication covering the entire Atari community. Reprint permission is granted, unless otherwise noted at the beginning of any article, to Atari user groups and not for profit publications only under the following terms: articles must remain unedited and include the issue number and author at the top of each article reprinted. Other reprints granted upon approval of request. Send requests to: dpj@delphi.com

No issue of Atari Online News, Etc. may be included on any commercial media, nor uploaded or transmitted to any commercial online service or internet site, in whole or in part, by any agent or means, without the expressed consent or permission from the Publisher or Editor of Atari Online News, Etc.

Opinions presented herein are those of the individual authors and do not necessarily reflect those of the staff, or of the publishers. All material herein is believed to be accurate at the time of publishing.